

Comment RC1:

- 1) This paper presents two independent analyses: minimization of the mass of a 15-MW floating platform and HRC steel price forecasting to 2033. Although the analyses are both connected to floating offshore wind LCOE, they are conducted separately and applied sequentially, which does not appear to result in any novel insights into potential systems interactions for economic and design optimization. Each analysis appears to be based on previously established methods. The abstract lists the best values from each analysis but does not present a combined result.
 - a. Dear reviewer, thank you for your comment. I agree with the assessment on the benefit of a more integrated approach to the technical and economic optimization. Both models were developed separately but a combined future analysis could include economic constraints within the design optimization framework. This could yield additional benefits especially if design constraints regarding motion are tightened, which could limit the potential for mass reductions, and a multi-objective framework could improve results. An integrated approach could allow for certain LCoE targets to be set which govern economic constraints and provide a geospatial context alongside auctions and could form an interesting development. Also, a more system-integrated approach could consider the platform dimensions and optimized reductions within the context of the mooring system, RNA and tower etc, alongside benefits for installation, port requirements and O&M processes.
 - b. What we did to correct:
 - i. Increase load cases
 - ii. Develop TSF model:
 1. Added structure to forecasts, T/T split, accuracy measures
 2. Selection of exogenous data?
 - iii. Combined results? Matrix of results across varying LCs?
 - iv. Adjustment to abstract showing revised results
 - c. In regards to the novel nature of the work, the optimization study here has nuance in its application to the Volturn-US platform using the developed design optimization framework, and applying its results to a detailed cost analysis and its integration within the RAFT framework, which solves the amplitude response of the floating turbine, under the design conditions and constraints at each iterative step of the design optimization. The economic method takes proven methods (SARIMAX long term, XGBoost short term) and applies them to structural steel which for marine renewables presents a novel approach in this area, alongside a hybrid forecast which selects the best performing model at each assessment interval. The time-series forecasting methods incorporated into the models were applied specifically to improve the predictions for materials with a direct effect on floating wind LCoE, which to the author's knowledge has not been done previously in literature in this context, and is of high relevance and importance considering current CAPEX and project financing concerns.
 - d. What we did to correct:
 - i. Better show the hybrid and predictive forecasts, the method behind feature selection

- ii. More insights into the exogenous data, their selection and impact onto the model
 - iii. Further show the improvements from the RAFT and GA optimisation over the mass reduction
- 2) The literature review includes several relevant citations for floating offshore wind platform design and cost assessment, however it lacks any references to previous academic work on modeling steel or commodity prices. This section could also be improved by explaining how each cited work is relevant to the current research rather than simply summarizing its contents.
 - a. I agree with the limitations of the literature review. Further citations could be included on structural steel forecasting, outside the marine renewables context which are available and can be included in a future iteration of this paper and would add more context to the study and could serve as a comparison to the accuracy and results of the timeseries forecasts.
 - b. What we did to fix
 - i. Added literature review of timeseries forecasting methods, and specifically those utilised in this presented research, namely SARIMAX and machine learning methods.
 - ii. Added literature for steel forecasting and. Linked to current research and this work and link
 - iii. Design optimisation literature also linked to our research and proposed paper
- 3) The methods section provides a clear description of the design optimization methodology. The price forecasting is less clear. Specifically, input data are discussed only generally and symbols in Eq. 6 are not defined.
 - a. We defined the method within Figure 2, but this was high-level and does not capture all information flows and processes which may help the reader understand the method, including the internal processes in the forecasting tools. This would also raise the level of detail to match Figure 2a. Equation (6) term were mistakenly omitted which is an easily fixed oversight. Section 3.3.2 could be developed further. Especially with LCoE calculations and the integration of generic CAPEX items without direct influence on this study's focus. There are more figures which could better explain the economic method but were omitted due to the WESC page limit, but can be rectified through text revision, and a more detailed process flow chart and explanation, and perhaps more explanatory process figures if space allows.
 - b. What we did to fix
 - i. We have restructured the section into 3.1 Model for platform optimisation and 3.2 Price forecasting and Economic assessment, and then two subsections within to add clarity.
 - ii. Equation 6 now defined with terms
 - iii. Within the new subsections the input data. Better explanation of the exogenous data selection and filter method.
 - iv. Improved explanation of the TSF method and model behaviour
- 4) Table 2 presents the results of the design optimization in a complete and concise format. Figure 4 is also a useful representation of the price forecasting results. Figure 5 appears to show poor performance of the forecast models (very high RMSE, MASE, and

sMAPE) and it is difficult to identify which axis (R or L) corresponds to which gray lines in the upper left subfigure.

- a. On Figure 5, (L) and (R) denote the relevant axis but this could be improved and I do agree it is hard to read the season-naïve and zero-mean forecasts for each. Careful selection here could also provide other clear figures which better show the model selection and performance, and will be considered for a revised edition, such as limiting the charts to RMSE or MASE, and including possibly timeseries forecast evolutions with different input criteria, model types or train-test splits.
- b. What we did to fix:
 - i. Add to Table 2 as more results and bounds can be improved upon
 - ii. Redid the figure analysing the models which now has clear comparisons of the three model types

We thank the reviewer for these insightful comments about our submission. A future revision can take on board this feedback and build a stronger literature section to understand how to make this novel contribution to floating offshore wind, which is through this optimization and future price forecasting tool. We believe time-series forecasting within the framework of floating offshore wind design is a new and exciting challenge, and a revised edition will aim to make this objective known with greater clarity. Kind regards.

Comment RC2:

- 1) The study addresses critical challenges currently threatening the industrial development of the floating offshore wind sector: **high project costs and limited visibility on raw material price trends**. It is structured around two main axes. First, a genetic algorithm combined with a numerical model is employed to optimize the design of a reference platform, in terms of mass reduction, under a set of constraints. Second, a machine-learning approach is introduced for short- to medium-term steel price forecasting, followed by an economic assessment of the optimized design. A conservative approach was adopted for selecting the optimum design, which still delivers **notable results (9% mass reduction)** and demonstrates potential for reducing FOW costs. The price forecasting component showcases an interesting combination of machine-learning models; however, the **promising outcome (2% LCOE reduction)** should be viewed in light of the **model's high uncertainty**. The work is presented with the scientific rigor required for publication in a referenced journal such as Wind Energy Science. Nevertheless, readers unfamiliar with machine-learning concepts may find it **challenging to grasp the subtleties** of the forecasting methodology and its evaluation. While the study does not introduce **significant scientific novelty**, relying on established methods for both design optimization and forecasting, **the integration of these approaches is new** and the results provide a promising foundation for further research. What the study lacks in methodological innovation, it compensates for with strong industrial relevance and potential for future cost reductions. **Overall, the manuscript is well-structured and clearly written.**
 - a. Dear reviewer, thank you for your comment. We agree that the promising LCoE reduction needs to be taken under the consideration of being within a prediction and that uncertainties must be well defined and explored through combining a multitude of design variables to lead to one or a group of forecasts, and that the

forecast range is not so long as to explode the uncertainty confidence bounds, which are shown in Fig. 4 in the 95% band. Regarding the technical description of the machine-learning aspects of the model, these can certainly be adjusted to improve the message and the findings of this work to a wider scientific audience, through clearer description of the equation terms and perhaps linking the forecasting steps better and with examples. Also, a page limit constraint prevented further figures depicting the development of the model and how the predictions were evaluated against real data. A review on this could certainly aid the reader's comprehension of the forecasting method.

- b. What we did to fix
 - i. Ensure price forecasting terminology is well defined and explained
 - ii. Explain and define the model uncertainty with results, add range of results and confidences to boost certainty
- 2) The key industrial issues mentioned above are **well introduced and clearly justify** the relevance of the work. **The literature review is extensive** and provides a detailed overview of historical research and the current state of the art in structural design optimization, floating platform numerical modeling, genetic algorithms, and FOW economics. **However, the scientific gaps in current practices**—particularly in design optimization and steel price forecasting—**could be more explicitly identified** to better situate the study within the existing research landscape.
 - a. We agree with this comment; the review perhaps leans more on the initial work in developing these fields and could be improved with a review into current advances in both time series forecasting and design optimization. Also, the applications of these sciences into current floating offshore wind or related marine energy operations.
 - b. What we did to fix
 - i. We have made improvements to the literature review to address this
 - ii. Gaps in scientific knowledge, mainly from the lack of combined design and economic analysis have been included where possible (e.g. Barter et al 2020” on areas of future work
 - iii. Clearer literature review on the time series forecasting methods and to address gaps in knowledge where possible.
- 3) The overall structure is clearly presented, and the quality of the literature review helps readers without prior expertise understand the approach. However, the **key quantities of interest derived from the numerical simulations—and used as constraints in the optimization problem—should be detailed in this section rather than in the results**, as they are essential indicators of the study's validity. Including references to relevant standards would also strongly support the study. **The Design Load Cases (DLC) against which the designs are tested should also be mentioned here**, instead of in the results section, for clarity.
 - a. variables which limit the design space of the optimization should be defined and backed up from past peer-reviewed work and placed within the method, rather than being defined and described within the results. For the key values derived from the numerical modelling, they were placed within the results section as the optimized values are shown alongside the original values and bounds. However, they should be introduced and defined before this section so the reader understands the design optimization and floater response analysis process, and this will be addressed by either a new table in the method section and or a

different figure to show the floater adjustment from the genetic algorithm. The constraint boundaries will also be explained with their expected effect on results.

- b. What we did to fix:
 - i. Better definition of load cases (LCs), not DLCs which are an incorrect definition here.
- c. For the DLCs, we agree these are defined and are not a result of this work, but a selected DLC from which metocean conditions are defined which govern the load spectra acting on the floating turbine, and these values serving as simulation inputs should be moved to the previous section.
- d. What we did to fix
 - i. Link to standard where possible, however this work is in the frequency domain and exact applications to DLCs (DNV / EIC) are hard to link
 - ii. Made adjustments throughout the manuscript to correct error stating “DLC” to ELC” or “LC” which is a more fitting term for the frequency domain load cases
 - iii. Table 1 now shows the varying load cases and their link to the reference DLCs where possible

4) The manuscript shows great attention to explaining how the algorithm and numerical model work, which is valuable for readers’ understanding. However, **the methodology lacks important details and overall clarity**. For instance, **equation (6) is not explained** and therefore appears of limited relevance; providing a detailed description of this formula would help readers understand the **functioning of the SARIMAX model**. Additionally, a clearer identification of the training and testing datasets, as well as the evaluation method used for the models, would significantly improve transparency. Overall, the **steel price forecasting section remains somewhat obscure for readers without machine-learning expertise**. Furthermore, steel prices are highly volatile and influenced by geopolitical and other exogenous factors; **incorporating scenario-based analysis** could strengthen the study. The link from steel price forecasting to LCOE is, however, clear and well justified.

- a. Equation 6 does indeed have missing the explanation of terms, this must have been an issue in final editing and transfer to PDF, and will be corrected for a final version. Regarding the training and testing datasets, these figures would indeed add an interesting element to the selection of the final model parameters and which train/test split yields the best results. Evaluation criteria for the model selection was using well-established methods such as mean error, mean average error, and model evaluation metrics such as (Akaike/Bayesian Information Criterion) were used and the outputs from these can be included. I do agree that more information on the training data and also the exogenous data that helps build the model predictions would add value and information to the reader.
- b. What we did to fix
 - i. Eq. 6 terms added and explanation.
 - ii. Added train/test hindcast plots
 - iii. Add clarity to the steel forecasting method
 - iv. Scenarios analysis has now been included by way model structure. Although hard to specify the exact scenario as model inputs are

timeseries data but is an expected future work. Pessimistic exog selection for example or more optimistic supporting data.

- 5) Some important details, such as the **numerical model outputs chosen as constraints for the design optimization and the DLC**, should be included in the methodology for greater clarity. **The constraint limits should also be described in more detail**, as they represent the most restrictive factors in the design optimization and therefore require precision. The design optimization results are solid and appear to provide a good basis for further research, as stated in the manuscript.
 - a. Already answered this question regarding DLCs
 - b. What we did to fix
 - i. Changed DLC to LC and listed in tables more information on LCs
- 6) **The evaluation of the ML model (supported by Figure 5) is difficult to interpret.** What are the baselines (seasonal naïve and zero-mean)? Against which reference are the different accuracy indicators calculated? These indicators, while probably familiar to anyone working in ML, should at least be named and defined for clarity.
 - a. Regarding the model evaluation, the results and assessment could certainly be clearer, especially for those without a background in machine learning. For the seasonal naïve, the baseline assumes that future prices assume the most recent seasonal pattern. For the zero-mean, the forecast is equal to a mean reference value and can be thought of as a lower boundary to define the performance of a model prediction against one with no information.
 - b. When comparing the models using the indicators, RMSE is computed for each model separately against the two aforementioned baselines in the top-left plot. The top right plot is the improvement (or reduction of errors) against the seasonal-naïve, we agree here with the following comment that the y-axis scales here are confusing and hard to interpret and will be scaled differently in the following iteration. For the MASE, values are against seasonal naïve for a range of future time-steps (or horizons), 1-12 here represents 1-12 months ahead. The sMAPE also looks to have an irregularity in the calculation as they should both cross (SARIMAX crossing at around horizon 6) in both plots. This will be addressed in the future iteration. We agree that these terms need further clarification for the reader to understand how these metrics are used to evaluate the models performance in forecasting.
 - c. For the DLCs, we agree these are defined and are not a result of this work, but a selected DLC from which metocean conditions are defined which govern the load spectra acting on the floating turbine, and these values serving as simulation inputs should be moved to the previous section.
 - d. What we did to fix:
 - i. These plots have now been adjust to more readable and informative replacements that show the three model's performance across both overall prediction success and per horizon forecast.
- 7) The LCOE results are promising (regarding the benefits of the mass redutcion granted by the design optimization) **but lack critical discussion given the very high variability in steel prices** (ranging from around €4,500 to €7,000 in a 5 years horizon). Interestingly, the XGB model is presented as the most accurate forecaster, **yet produces the most pessimistic scenario, which is unusual**. This raises questions about the **relevance of the hybrid method** if it is not the most reliable option compared to short-term-oriented

XGB and long-term-oriented SARIMAX. Again, all of this might be clear to readers well-versed in ML, but less so for others.

- a. We agree with the comment on a lack of critical discussion, especially considering this work has high industrial relevance especially in the current economic environment for floating offshore wind. Results will be placed in context of auctions, active projects etc to add real world comparisons to these research outcomes. Based on these evaluation metrics, the XGB forecast does indeed yield the most accurate results within the testing window, and does lead to the most pessimistic forecast. This outcome was combined with the SARIMAX in a hybrid to try and accommodate the best of both models for a single forecast. In hindsight, perhaps more single model outputs using varying criteria (for example, different sets of train/test splits or exogenous data groups) would be useful. This would also confirm the previous comment that a scenario approach could boost the results, and will be taken into consideration.
 - b. What we did to fix:
 - i. After model adjustment, the hybrid model now performs best which is now detailed in the method and results
 - ii. Better method now describes the forecasting and ML methods to assist with comprehension of these topics
- 8) The choice of figures is generally very good and supports the reader's understanding. However, **Figure 5 could benefit from more explanation in the text**, as noted above. In addition, **the y-axis of some graphs appears unusual** (e.g., RMSE of 0.050 €/ton of steel; -500,000% gain). Providing additional clarification would help readers interpret these values correctly.
- a. What we did to fix:
 - i. These plots have now been redone with new figures and economic analysis
- 9) Some approximative phrasing affects the reader's comprehension. **A non-exhaustive list includes lines 13 and 30**. Additionally, **Equation (3) should be rewritten for better clarity**, as its current form can be interpreted as placing the second moment of area in the numerator.
- a. On the approximate phrasing, yes, those lines yield confusion especially in the abstract where results are defined without full explanation, which will be addressed as the abstract is often the first point of information about this work and enhances the papers' reach. On the equation, yes it is ambiguous relating to the denominator and will be adjusted.
 - b. What we did to fix
 - i. The abstract has now been adjust to reflect new results
 - ii. Added parenthesis to Eq3 for clarity.

My response (sent 15.01.25):

Dear reviewer, thank you for your comment. We agree that the promising LCoE reduction needs to be taken under the consideration of being within a prediction and that uncertainties must be well defined and explored through combining a multitude of design variables to lead to one or a group of forecasts, and that the forecast range is not so long as to explode the uncertainty

confidence bounds, which are shown in Fig. 4 in the 95% band. Regarding the technical description of the machine-learning aspects of the model, these can certainly be adjusted to improve the message and the findings of this work to a wider scientific audience, through clearer description of the equation terms and perhaps linking the forecasting steps better and with examples. Also, a page limit constraint prevented further figures depicting the development of the model and how the predictions were evaluated against real data. A review on this could certainly aid the reader's comprehension of the forecasting method.

We agree with this comment; the review perhaps leans more on the initial work in developing these fields and could be improved with a review into current advances in both time series forecasting and design optimization. Also, the applications of these sciences into current floating offshore wind or related marine energy operations.

The manuscript can be reformatted to accommodate these suggestions and they are valid. Key variables which limit the design space of the optimization should be defined and backed up from past peer-reviewed work and placed within the method, rather than being defined and described within the results. For the key values derived from the numerical modelling, they were placed within the results section as the optimized values are shown alongside the original values and bounds. However, they should be introduced and defined before this section so the reader understands the design optimization and floater response analysis process, and this will be addressed by either a new table in the method section and or a different figure to show the floater adjustment from the genetic algorithm. The constraint boundaries will also be explained with their expected effect on results.

For the DLCs, we agree these are defined and are not a result of this work, but a selected DLC from which metocean conditions are defined which govern the load spectra acting on the floating turbine, and these values serving as simulation inputs should be moved to the previous section.

Equation 6 does indeed have missing the explanation of terms, this must have been an issue in final editing and transfer to PDF and will be corrected for a final version. Regarding the training and testing datasets, these figures would indeed add an interesting element to the selection of the final model parameters and which train/test split yields the best results. Evaluation criteria for the model selection were using well-established methods such as mean error, mean average error, and model evaluation metrics such as (Akaike/Bayesian Information Criterion) were used and the outputs from these can be included. I do agree that more information on the training data and also the exogenous data that helps build the model predictions would add value and information to the reader.

Regarding the model evaluation, the results and assessment could certainly be clearer, especially for those without a background in machine learning. For the seasonal naïve, the baseline assumes that future prices assume the most recent seasonal pattern. For the zero-mean, the forecast is equal to a mean reference value and can be thought of as a lower boundary to define the performance of a model prediction against one with no information.

When comparing the models using the indicators, RMSE is computed for each model separately against the two aforementioned baselines in the top-left plot. The top right plot is the improvement (or reduction of errors) against the seasonal-naïve, we agree here with the following comment that the y-axis scales here are confusing and hard to interpret and will be scaled differently in the following iteration. For the MASE, values are against seasonal naïve for

a range of future time-steps (or horizons), 1-12 here represents 1-12 months ahead. The sMAPE also looks to have an irregularity in the calculation as they should both cross (SARIMAX crossing at around horizon 6) in both plots. This will be addressed in the future iteration. We agree that these terms need further clarification for the reader to understand how these metrics are used to evaluate the model's performance in forecasting.

We agree with the comment on a lack of critical discussion, especially considering this work has high industrial relevance especially in the current economic environment for floating offshore wind. Results will be placed in context of auctions, active projects etc to add real world comparisons to these research outcomes. Based on these evaluation metrics, the XGB forecast does indeed yield the most accurate results within the testing window and does lead to the most pessimistic forecast. This outcome was combined with the SARIMAX in a hybrid to try and accommodate the best of both models for a single forecast. In hindsight, perhaps more single model outputs using varying criteria (for example, different sets of train/test splits or exogenous data groups) would be useful. This would also confirm the previous comment that a scenario approach could boost the results and will be taken into consideration.

On the approximate phrasing, yes, those lines yield confusion especially in the abstract where results are defined without full explanation, which will be addressed as the abstract is often the first point of information about this work and enhances the papers' reach. On the equation, yes, it is ambiguous relating to the denominator and will be adjusted.