Dear reviewer,

Thank you for your time and effort spent into the review of our manuscript. We hope this reply sufficiently answers your questions.

1. Penetration depth, defect detectability, and layer-thickness extraction are treated mainly qualitatively. Quantitative rigor is insufficient in several key analyses. The study would benefit from statistical summaries over multiple locations and more formal error metrics.

Thanks for your comment. For the maximum penetration depth, this was calculated through the comparison of the noise floor of our OCT system and the measurement of the sample outside any intentional defect. The noise floor was measured when there is no sample under the head scan, as presented in figures 8 and 9. Penetration depth is intrinsically dependent on the system combined with the chosen sample, representation one layer composition example applied within the wind turbine blade industry.

In addition, as all OCT B-scans contain speckle noise, it was necessary to compare our measurements with the noise floor of our system to ensure rigorous validation to distinguish between scattering from the material and sample enabled speckle patterns.

2. The attenuation-based interface-detection algorithm seems underdeveloped. The conceptual description (Fig. 5) is helpful, but the lack of methodological detail, no validation against ground truth, and poor performance in defected areas weaken the Conclusions. Suggested claims about future integration of deep learning are reasonable, but they should not compensate for the current missing quantitative evaluation.

About the algorithm, we agree; this algorithm is very basic and follows the attenuation slopes of the different layers.

The validation slope parameters (*objective fits*) have been measured outside the intentional defect area with a large quantity of data averaged to ensure the slope value extracted was not influenced by any noise or manufacturing defect. The data size represents 100 B-scans composed of 1000 A-scans each, so the slope calculation for the objective fits is based on 10^5 mean A-scans.

Then these *objective fits* have been used in the algorithm to find the layer over the B-scan studied.

This algorithm has been made to help us to isolate the different layers more easily, but as it is a primitive version, we still see room for improvement, e.g. by increasing the dataset size. About the deep learning integration, this one is already in progress with our project partner. We would like to implement deep learning solutions to automatise the layer thickness detection process during manufacturing. Combining larger datasets and knowledge transfer, we are working on improving the precision of coating layer interface delineation.

3. Industrial scalability claims seem overstated. The speculative calculation about faster scanning times lacks a realistic time-budget analysis and seems optimistic. These observations should be presented more cautiously, experimentally backed up, or omitted.

Thanks to the reviewer for this comment. Yes, indeed, the scalability seems to be not feasible with the speed value given from the experimental one. We dare extrapolate on our results as we are collaborating with NORBLIS (https://norblis.com/)developing industrial MIR OCT scanners suited for industrial demands. A LinkedIn post has been published for a milestone on this development (https://shorturl.at/VQJOw).

4. Key procedures, including refractive-index derivation, OCT post-processing, and sample preparation, need clearer and more rigorous descriptions to allow replicability.

The refractive indexes have been calculated through the conventional method as explained on page 8. Using the sample with aluminium foil inside. Optical path delay (OPD) in the air represents the physical distance, as the refractive index of the air is approximately equal to 1 for a temperature of 21 degrees.

We extracted the physical distance $L_{topcoat}$ and $L_{topcoat+primer}$ by removing the layer above the corresponding aluminium foil and measured it through our NIR OCT B-scans. Then, as it was possible to identify these two aluminium foils with the MIR OCT system, we calculated the corresponding refractive index.

The OCT post-processing is explained in this article https://doi.org/10.1038/s41377-019-0122-5 as presented inside the article in material and method.

About sample preparation, they were made by Siemens Gamesa, and we are not allowed to share the detailed procedure of it.

5. Related to the above, the description of the sample preparation requires greater detail.

This is a good suggestion, but unfortunately, as these samples were made by Siemens Gamesa we cannot provide many more details.

6. The state-of-the-art for damage assessment in GFRP is relatively limited. Many recent contributions can be found for this specific application, for example, in the field of vibration-based monitoring, https://doi.org/10.1016/j.compstruct.2020.112882 and https://doi.org/10.1002/stc.2805. These could be added to the discussion in the Introduction.

Thanks for your notice, as the paper focuses on the characterisation of the coating above the GRFP of the wind turbine blade (WTB) because the state-of-the-art for damage assessment in GRFP is poor. The idea here was mainly to provide a quick representation of how the different non-destructive techniques are used in WTB. But we will cite and include the most recent article of the two proposed.

7. The text repeatedly highlights penetration of "~360 µm". However, it seems that this value corresponds to a single dataset (sample no. 1) and thus is not a generally valid statement. It is one value of one measurement. The authors should provide statistics across several scans.

For the maximum penetration depth, this one was calculated through the comparison of the noise floor of our OCT system and the measurement of the sample outside any intentional defect. This value is the maximum that we could obtain. For other similar samples we screened, the penetration with slightly lower but beyond 300 µm .

8. The figures could benefit from higher contrast and more explicit annotation

Thanks for the suggestion, but the figures are already very saturated by the annotation. By adding more, we are afraid that we will compromise the reading and the comprehension of the manuscript.

9. The English is overall good, but there are occasional grammatical errors, such as missing articles. A thorough review could be useful.

Thanks for noticing this; the latest typos have been corrected.