Dear reviewers,

We thank both reviewers for their careful reading of the manuscript and for the constructive comments and suggestions, which have substantially improved the clarity, precision, and completeness of the paper. Below we respond to each comment individually and detail the corresponding revisions made to the manuscript.

REFEREE COMMENT 1 (RC1)

Comment 1 – Experimental setup

Without breaking confidentiality, would it be possible to provide more information on the wind farm from where the data was colected? Additionally, would it be possible to explore the effects of wakes on the turbine loads, if any?

We have added a description of the wind farm to Section 3, characterising it as a multi-megawatt onshore installation located along a mountain ridge crest in central Portugal, with turbines arranged in an elongated, quasi-linear layout over karstic, highly irregular terrain that induces non-uniform wind flow and locally enhanced turbulence.

Regarding wake effects, we note that the present study is based on strain measurements from a single instrumented turbine. Without simultaneous measurements from neighbouring turbines, a quantitative, turbine-resolved assessment of wake effects on tower fatigue is not possible with the current dataset. We have nonetheless added a qualitative note acknowledging the ridge-crest layout as a source of potential wake interactions, and we identify a direction-resolved analysis of turbulence and wake effects as a valuable direction for future work, requiring simultaneous strain measurements from multiple turbines across the array.

Proposed manuscript revision:

- Section 3, first paragraph inserted:
  "The wind farm is a multi-megawatt onshore installation located along a mountain ridge crest in central Portugal. The wind turbines are arranged in an elongated, quasi-linear layout along the ridge, over karstic, highly irregular terrain that, combined with heterogeneous land cover, induces non-uniform wind flow and locally enhanced turbulence."
- Section 3, inserted after the SHM/instrumentation paragraph:
  "Given the ridge-crest, quasi-linear turbine layout, wake interactions with neighbouring turbines are expected for wind directions aligned with the ridge axis, potentially increasing unsteady thrust loading. As this study is based on strain measurements from a single instrumented turbine, a quantitative assessment of wake effects on tower fatigue was beyond the scope of the present dataset."
- Section 5, inserted after the SHM/instrumentation paragraph:
  "Further valuable directions for future work include: a direction-resolved analysis of turbulence and possible wake effects on tower fatigue, requiring simultaneous strain measurements from multiple turbines across the array."

Comment 2 – Extrapolation

Would it be possible to include an additional study investigating whether the fatigue damage estimated at the instrumented tower section can be extrapolated to other locations along the tower, such as the tower top or tower base? These locations are often of greater practical interest for structural health monitoring and lifetime assessment. Even if such an analysis is beyond the scope of the present work, a discussion of the challenges and limitations associated with extrapolating damage estimates to other structural locations would strengthen the paper.

The fatigue estimates presented in this study are obtained directly from strain measurements at the instrumented sections, T1 and T2, and can therefore be validated against measured data, what we would describe as reconstruction of fatigue damage. Extrapolating these estimates to other locations along the tower, such as the base or top, is a distinct problem: it cannot be validated directly against measurements at those (non instrumented) locations and instead requires an explicit load-transfer relationship between the instrumented sections and the location of interest — for example, derived from a structural or modal model, or from a data-driven prediction approach. Establishing and validating such a relationship is outside the scope of the present dataset, since it would require either additional instrumentation at the target locations for validation, or a structural model of the tower that has itself been validated against measurements.

We agree, however, that discussing this distinction and its associated challenges strengthens the paper, and we have added text addressing it, together with a note identifying it as a direction for future work.

Proposed manuscript revision:

- Section 4, at the end of the discussion of results:
  "The fatigue damage estimates presented above relate specifically to the instrumented tower sections, T1 and T2. Extrapolating these estimates to other locations of practical interest, such as the tower base or tower top, is a distinct problem from the reconstruction addressed here: it cannot be validated directly against measurements at those locations and instead requires an explicit load-transfer relationship, for example derived from a structural/modal model or from a data-driven prediction approach, established between the instrumented sections and the location of interest. Differences in local dynamic response, modal contributions and load-path characteristics along the tower height mean that such a relationship cannot be assumed to hold without dedicated validation."
- Section 5, added to the future-work paragraph:
  "Further valuable directions for future work include: (…) extrapolation of fatigue damage estimates from the instrumented tower sections to other locations of interest, such as the tower base or tower top, through validated load-transfer or data-driven relationships."

Comment 3 – Discussion on the damage calculation

For completeness, the authors could provide additional details on the fatigue damage calculation methodologies considered in the study. Including the governing equations or a brief description of the underlying algorithms for each approach would improve the paper's accessibility and allow readers to better understand the differences between the methods.

The damage-accumulation formulation (Palmgren–Miner rule, Equation 1, and the EC3/DNV σ-N curves) is already presented in Section 2.5, lines 228–237. What was missing was an explicit, unified statement of how the cycle spectrum (n_i, $\Delta\sigma_i$) that feeds into Equation (1) is obtained by each of the methods compared in Section 4, this is precisely what distinguishes the methods from one another. We have added a short paragraph immediately after line 238 (end of the $\sigma - N$ curve paragraph) and before line 239 ("The specific metrics adopted in this work..."), so that the reader has this description in hand before reaching Table 1 and the results in Section 4.

Proposed manuscript revision:

- Section 2.5, new paragraph inserted between line 238 and line 239:
  "All fatigue-monitoring approaches compared in Section 4 apply the damage formulation of Equation (1) to a cycle spectrum (n_i, $\Delta\sigma_i$) and differ only in how this spectrum is obtained from the measured stress signal. Continuous RFC applies rainflow counting directly to the full stress time series and is adopted as the baseline. Window-based RFC (Section 2.2) applies rainflow counting independently within each aggregation window, discarding cycles whose turning points span window boundaries. LFFD (Section 2.3, Figure 2) recovers these boundary-spanning cycles by additionally applying RFC to the concatenated sequence of unclosed half-cycles (residuals) left at each window boundary, combining this low-frequency contribution (L_{f+h}) with the within-window, high-frequency contribution (H_{f+h}) to recover the continuous-RFC spectrum exactly. SPVE and Max–Min (Section 2.4, Figure 3) instead replace the full stress time series within each window by a small set of representative extrema, start-peak-valley-end, or maximum/minimum, respectively, before a single RFC pass is applied to the concatenated extrema sequence."

Comment 4 – Operational States

Did the authors consider analyzing the impact of different operational states (e.g., start-up, shutdown, normal operation, and idling) on fatigue damage accumulation? If not, could they comment on why this analysis was omitted? Understanding the contribution of individual operational states, particularly their role in generating low-frequency, high-amplitude load cycles, would provide valuable insights into the mechanisms driving fatigue damage and further strengthen the paper.

As reported in Section 3 (lines 282–285), the analysed six-month period is dominated by the Run state (87.24% of samples), with Ready accounting for a further 11.49%; start-up, shutdown, idling and other transient states together represent a small fraction of the dataset. A statistically robust, per-state breakdown of the fatigue-damage contribution was not pursued in the present study, since the low relative occurrence of transient states would require a substantially longer monitoring record to obtain damage estimates for each state with adequate confidence.

We agree, however, that operational-state transitions are a physically relevant source of low-frequency, high-amplitude cycles, and that this connects directly to the central theme of the paper. Such transients (e.g. impulsive pitch actuation during start-up/shutdown, or the free-yawing/feathered-blade response during idling) are precisely the kind of long-period, boundary-spanning events that conventional window-based rainflow counting under-represents, and which the low-frequency (residual) component of LFFD is designed to recover. They are therefore already implicitly captured, rather than isolated, within the low-frequency damage contribution reported in Section 4. We have added a paragraph clarifying this link and identifying a quantitative, per-state damage decomposition as valuable future work.

<u>Proposed manuscript revision:</u>

- <u>Section 4.4, inserted between line 417 and line 418:</u>
  "This link between operational-state transitions and the low-frequency damage contribution motivates a more detailed characterisation of individual states (e.g. start-up, shutdown, idling). However, the analysed period is dominated by steady production states (Run: 87.24%; Ready: 11.49%; Section 3), with transient states occurring comparatively rarely; a statistically robust, per-state decomposition of fatigue damage was therefore not pursued in this study and would require a longer monitoring dataset. It is identified as a valuable direction for future work (Section 5)."
- <u>Section 5, added to the future-work paragraph, alongside the wake/turbulence and extrapolation items already added:</u>
  "Further valuable directions for future work include: (…) a quantitative decomposition of fatigue damage by operational state, isolating the contribution of start-up, shutdown and idling transients from steady production."

Comment 5 – Cluttered figures

Some of the figures in Section 4, particularly Figures 8, 10, and 11, appear visually cluttered and contain a large amount of information. I recommend simplifying these figures or splitting them into multiple panels to improve readability and make the key findings easier to interpret.

We reviewed each of the three figures individually and propose the following simplifications:

- **Figure 8** (cumulative FA-direction damage at T1, panels a–b for 10-minute and 15-second windows) currently overlays four sampling rates (50, 25, 12.5, 6.25 Hz) × three damage components (total, low-frequency, high-frequency) in each panel, i.e. up to 12 curves per panel, with end-of-period annotations. We propose splitting each panel into three stacked sub-panels, one per damage component (total / low-frequency / high-frequency), each showing the four sampling rates. This preserves all reported values while removing the overlap between components.

- **Figure 10** (FA-direction fatigue at T1, panels a–d, comparing 10-minute vs 15-second windows and EC3 vs DNV) currently shows total/low/high-frequency curves for both $\sigma - N$ standards within the same panels. We propose separating the EC3 and DNV results into two figures (or a 2×4 panel grid clearly separated by row).
- **Figure 11** (comparison of LFFD, window-based counting, SPVE and Max–Min, showing cumulative damage and $\Delta\sigma - N/\Delta\sigma - D$ spectra) currently overlays all four methods across both window lengths in the same panels. We propose splitting this into two figures: one showing cumulative damage curves (window-length comparison), and a second showing the end-of-period $\Delta\sigma - N$ and $\Delta\sigma - D$ spectra, each with a simplified legend grouping methods by colour and window length by line style.

Across all three figures, we will also standardise the colour/line-style convention (e.g. colour for damage component or method, line style for window length or $\sigma - N$ standard) so that the same visual code is reused consistently, reducing the cognitive load of interpreting each figure.

Proposed manuscript revision:

- Figures 8, 10, 11 and 12 (renumbered from the original submission) were reorganised into multi-panel layouts:
  Figure 8 now shows one sub-panel per damage component with a zoomed inset of the final month; Figures 10 and 11 separate the EC3 and DNV results into independent figures; and Figures 12 and 13 separate the five compared fatigue-monitoring methods into two thematically grouped figures (residual-preserving vs. window-based; LFFD vs. reduced-information extrema representations).

REFEREE COMMENT 2 (RC2)

Comment 1

Line 13: It is very rare to install sensors on all turbines of a farm. So, this claim that high-frequency data is impractical at fleet scale is not relevant.

We agree the original wording is imprecise. Our intended point was not that every turbine in a fleet would need to be instrumented with strain sensors, but rather two related practical constraints: (i) even for the limited number of turbines that are instrumented, continuously storing and processing high-frequency strain data via rainflow counting over long-term operation is costly and rarely sustained in practice; and (ii) since strain instrumentation is only installed on a small subset of turbines to begin with, continuous-RFC-based fatigue assessment is simply unavailable for the vast majority of a fleet, which is why fatigue monitoring approaches usually rely on signals that are available fleet-wide, such as SCADA. We have revised the sentence in the Abstract to state this more precisely.

Proposed manuscript revision:

> "However, only a small subset of turbines within a fleet are typically instrumented with strain sensors, and even for these, the data volumes, computational complexity, and cost of continuously processing high-frequency strain data via rainflow counting over long-term operation make this approach impractical to sustain at scale. For this reason, industry practice has long relied on stress data aggregated over 10-minute windows already used by Supervisory Control and Data Acquisition (SCADA) systems, which are available across the full fleet."

Comment 2

Line 14 and 15: "This choice is convenient but tends to underestimate fatigue damage generated by cycles whose period exceeds the windows". We already know that by a simple cyclecount on the residuals signal, this limit is solved. So this sentence is not 100% correct.

We agree the sentence as worded is imprecise. It refers specifically to *conventional* window-based rainflow counting, which discards or truncates cycles spanning window boundaries, not to residual-preserving approaches, which, as the reviewer correctly notes, already recover this loss by counting cycles on the concatenated residual sequence. Rather than only qualifying the limitation, we have revised the abstract to state plainly that residual-preserving methods such as Low-Frequency Fatigue Dynamics (LFFD) already exist to address it, positioning the paper's contribution correctly: not the reconstruction of boundary-spanning cycles itself, but its systematic application and validation at sub-minute, SCADA-aligned window lengths.

Proposed manuscript revision:

> “This choice is convenient but, when combined with conventional window-based rainflow counting, tends to underestimate fatigue damage generated by cycles whose period exceeds the window length, driven by factors such as slow variations in wind speed and direction, operational transitions, control actions, and seasonal or daily temperature variation, as well as, for offshore turbines, changes in wave parameters. Residual-preserving approaches, most notably Low-Frequency Fatigue Dynamics (LFFD), address this limitation by reconstructing cycles that span window boundaries, but have so far been applied almost exclusively within the conventional 10-minute framework.”

Comments 3 and 5

Having sub-minute windows has higher SCADA resolution compared to 10-minute windows, as usually the statistical values for SCADA are used to calculate bending moments, or to define the operational state of each window. I missed this discussion in your abstract.

Line 19: Which cycles are discarded? We usually cycle count the high-frequency stress time-series within 10-min, so no cycles are discarded. The sentence needs clarification.

Regarding the first: the original abstract did not include this point because its scope was deliberately centred on the impact of sub-minute SCADA resolution specifically on fatigue damage computation, rather than on a broader characterisation of what changes between 10-minute and sub-minute SCADA data. We agree, however, that briefly acknowledging this additional application strengthens the motivation for adopting sub-minute resolution in the first place and have added a short clause to the Abstract accordingly, without expanding it into a broader discussion outside the paper's scope.

Regarding the second: we agree the wording "discards" is imprecise. Under conventional window-based rainflow counting applied to the full high-frequency stress time series, no cycle is literally discarded: every stress reversal is processed. What happens instead is that cycles whose turning points span a window boundary are truncated, split into two unclosed half-cycles (residuals) at that boundary, and, unless these residuals are subsequently reconstructed (e.g. via LFFD-style residual concatenation), the corresponding long-period, often large-amplitude cycles are never recovered as full cycles, and their damage contribution is systematically under-represented rather than dropped outright. We have revised the wording to reflect this precisely.

Proposed manuscript revision:

> “Technology has recently enabled operators to deploy second-resolution SCADA metrics at scale and at reasonable cost, offering finer-grained statistical descriptors for applications such as load and bending-moment estimation and operational-state classification. This is also an opportunity to revisit fatigue monitoring at a SCADA-aligned timescale that is short enough to limit the truncation of boundary-spanning cycles inherent to conventional window-based counting, yet aggregated enough to remain deployable across large fleets.”

Comment 4

Line 16: These parameters are not the only reasons behind low frequency cycles. It can also be wind direction, change in wave parameters in case of offshore turbines. It can even be the seasonal or daily temperature variation.

We agree the original list was not exhaustive and could be read as claiming these are the only drivers of low-frequency cycles. We have extended the phrasing to make clear this is a non-exhaustive, illustrative list, and added wind direction and seasonal/daily temperature variation as additional onshore-relevant examples; we also note wave-parameter variation as a further driver relevant to offshore turbines specifically, since this sentence is a general statement about the field and not limited to the onshore case study presented here.

Proposed manuscript revision:

> "This choice is convenient but, when combined with conventional window-based rainflow counting, tends to underestimate fatigue damage generated by cycles whose period exceeds the window length, driven by factors such as slow variations in wind speed and direction, operational transitions, control actions, and seasonal or daily temperature variation, as well as, for offshore turbines, changes in wave parameters."

Comment 6

Line 44: "based on" is not a correct wording. It is usually based on strain data. I guess you mean data discretization in 10-minute windows follows the 10-minute SCADA format. The 10-minute windowing is not just because of SCADA, but also due to how simulations are done in the design.

We thank the reviewer for this clarification and agree the original wording was imprecise on two counts. First, fatigue monitoring is fundamentally based on strain (or strain-derived stress) measurements, not on SCADA data itself; what follows the SCADA convention is the discretization of that stress signal into 10-minute windows, not the underlying measurement. Second, we agree that the 10-minute convention does not originate from SCADA systems alone: it is also embedded in how design-basis load simulations are performed and standardised (e.g. IEC 61400-1 load cases), which SCADA aggregation later came to mirror operationally. We have revised the sentence to reflect both points.

Proposed manuscript revision:

> "In practice, fatigue monitoring relies on strain measurements that are discretized into conventional 10-minute windows, following both design-basis load simulation practice and the aggregation format of Supervisory Control and Data Acquisition (SCADA) systems (IEC 61400-13, 2015)."

Comment 7

Line 49: What do you mean by aggregation over 10-minute windows? Similar comment as above. If it is high-frequency strain measurement, the high-frequency cycles should be covered. Did you mean that by cutting data into 10-minute windows and using the statistical SCADA parameters, the high-frequency cycles might be smoothed?

Yes, that is exactly what we intended, and we thank the reviewer for identifying the ambiguity. This sentence refers to the common fleet-wide practice of using SCADA statistical descriptors (means, standard deviations, extrema) as a proxy for fatigue-relevant loading, rather than cycle-counting the full strain signal directly, as the reviewer notes, when the full-resolution signal is itself rainflow-counted within 10-minute segments, no smoothing of that kind occurs (this is the boundary-truncation effect already clarified in our response to the earlier comment on Line 19). We have revised the sentence to make this distinction explicit.

Proposed manuscript revision:

> "When these 10-minute statistical descriptors, rather than the full-resolution strain signal, are used as a proxy for fatigue-relevant loading, this aggregation may therefore smooth short-term stress fluctuations and distort the underlying cycle distribution, potentially biasing fatigue estimates."

Comment 8

Line 66: Can you explain here shortly why should someone need this separation?

The explicit separation is needed because, as shown later in this paper, at short window lengths the large majority of fatigue damage is concentrated in the low-frequency, residual (inter-window) contribution, while the intra-window, high-frequency contribution becomes comparatively small. This has a direct practical motivation for this study: if the residual sequence alone can account for most of the fatigue damage, it becomes a natural, lightweight target for future prediction efforts, for example, predicting the residual sequence directly from SCADA signals, rather than requiring the full high-frequency strain signal to be transmitted and stored. Explicitly separating high- and low-frequency contributions is therefore the necessary first step towards that objective, since it isolates which component actually needs to be predicted or preserved. We have added a short clarification at first mention.

Proposed manuscript revision:

> "This explicit separation enables a more granular characterisation of fatigue behaviour and, as shown later in this study, is particularly relevant when the window length is reduced below the conventional 10-minute duration, since the low-frequency residual contribution then accounts for the majority of total fatigue damage."

Comment 9

Line 68: Can you discuss the limitations of reduced information strategies? For a general SHM campaign, can we have the FFT, PSD, strain calibration, virtual sensor signal? How accurate would be the 15s damage? Although longterm accumulated damage is important, many prediction models are based on 10min window. For example, in Figure 8, the green lines for 15s windows are very low, meaning that the 15s damage is not a good representative by itself. So if a machine learning model is trained based on the 15s data, it needs an extra information about low frequency cycles. While this problem is less significant for the 10minute data.

We agree this is an important limitation to state explicitly, though we are careful not to overstate it: the reduced-information representations evaluated here (SPVE, Max–Min) were designed and evaluated specifically for fatigue-damage estimation, preserving only the extrema needed for rainflow-equivalent cycle counting. Their suitability for other SHM applications, such as frequency-domain analyses (FFT, PSD) for modal identification or vibration-based damage detection, strain-gauge calibration checks, or virtual-sensing algorithms, was outside the scope of this study and has not been evaluated here; we therefore do not claim they are insufficient for these purposes, only that this remains an open question requiring dedicated investigation.

Proposed manuscript revision:

- Line 543:
  "The reduced-information representations evaluated here were designed and assessed specifically for fatigue-damage estimation; their suitability for other SHM applications, such as frequency-domain analyses (FFT, PSD) for modal identification or vibration-based damage detection, strain-gauge calibration checks, or virtual-sensing algorithms, falls outside the scope of this study and remains an open question for future work."

- Line 545:
  "(…) and by predicting the low-frequency residual sequence directly, which this study shows accounts for the large majority of fatigue damage at short window lengths while representing only ~7% of the raw data volume."

Comment 10

Line 144: 2.2 title should be High-frequency-cycle counting. High cycle is not defined in the manuscript.

"High-cycle" risks being confused with the standard fatigue-engineering term "high-cycle fatigue" (HCF), which is a different concept (a fatigue regime, not a cycle-counting method) and is not what we intend here. The paper consistently uses "high-frequency" versus "low-frequency" to distinguish intra-window from inter-window contributions elsewhere in the manuscript, so renaming this method label to match that terminology removes the ambiguity. We have renamed it consistently everywhere it appears as a method label: the Section 2.2 heading, and its use as a row label in Table 2 (literature comparison) and Table 3 (storage/computational comparison).

Proposed manuscript revision:

Location	Before	After
Section 2.2 heading	"Window-based High-cycle Counting"	"Window-based High-frequency-cycle Counting"
Table 2 (row, method attributed to Sadeghi et al., 2022)	"Window-based High-cycle Counting"	"Window-based High-frequency-cycle Counting"
Table 3 (row, computational effort)	"Window-based High-cycle Counting"	"Window-based High-frequency-cycle Counting"

Comment 11

In Figure 1, section € of figure, it should be sigma H f+h i. Also i missing in section g too. Usually the counted cycles are shown by small n. Capital N is for the maximum number of cycles that yields to failure.

We thank the reviewer for these two observations and have corrected both throughout Figure 1.

For the aggregated panels (e) and (g), the summation index was indeed missing, as the reviewer notes; these have been revised to $\sum \Delta\sigma_{H_{f+h,i}}$ and $\sum \Delta\sigma_{H_{f,i}}$, respectively, making explicit that each aggregated histogram sums the per-window contributions $H_{f+h,i}$ and $H_{f,i}$ across all windows i.

Regarding cycle-count notation: we agree with the reviewer. All histogram panels in Figure 1 (b, d, e, f, g, h, k, l) previously labelled the x-axis as "N" (number of cycles to failure, per Equation 1), when they in fact display counted cycles, which should be denoted with lowercase "n". This has been corrected throughout the figure.

Proposed manuscript revision:

Panel	Before	After
(e)	$\Delta\sigma$	$\sum \Delta\sigma_{H_{f+h,i}}$
(g)	$\Delta\sigma$	$\sum \Delta\sigma_{H_{f,i}}$
All histogram x-axes (b, d, e, f, g, h, k, l)	N	n

Comment 12

In Figure 2, the link between merged full-cycle histogram and the gray area should be a bit higher. Now it gives the impression that pink and green boxes are linked. There is also another link between SHM raw data 50Hz and the pink box which might not be necessary.

Regarding the first point: we agree the connecting line from "Damage histogram corresponding to the merged full cycles" (green box) to "Merge high- and low-frequency contributions" (grey box) was routed too low, passing close to the pink (low-frequency) block and visually suggesting a direct link between the green and pink boxes, which does not reflect the actual data flow. We have raised this connecting line so it runs clearly above the pink block.

Regarding the second point: this link is intentional and necessary rather than redundant. Extracting the low-frequency residual sequence requires access to the original raw stress values, in their correct temporal order, for the unclosed excursions remaining after full-cycle extraction, information that is not available from the half-cycle counts alone, which is all that standard per-window Rainflow Cycle Counting output provides. Rather than removing the link, we have added a clarifying sentence to the figure caption explaining this dependency.

Proposed manuscript revision:

- Raise the connecting line between the green box and the grey "Merge high- and low-frequency contributions" box, so it clearly passes above the pink box.
- Extend the caption of Figure 2 with the following clarification:
  "The link from the raw SHM data to the residual-concatenation step reflects the need for the original, temporally ordered stress values of the unclosed excursions, which are not retained in the half-cycle counts alone."

Comment 13

Figure 3, title of the orange box should be min-max values per window? Caption is wrong (elements d to g)

Regarding the orange box title: this is a copy-paste error. The title currently reads "Start–Peak–Valley–End stress sequences per window" (carried over from the yellow SPVE box above), when the orange box represents the Max–Min branch and should read "Maximum–Minimum values per window".

Regarding the caption: we agree it does not correctly correspond to the figure, which contains six labelled panels (a–f), not eight (a–h) as the original caption describes. The caption incorrectly assumed a separate panel showing the raw segmented time series before the SPVE extraction step, and separate panels for the RFC operation and its resulting spectrum, none of which are shown as distinct panels in the actual figure. We have rewritten the caption to match the six panels actually present.

In addition, and consistent with our response to the comment on Figure 1, we identified the same "N" versus "n" inconsistency in Figure 3: the x-axes of the resulting fatigue-spectrum panels (SPVE and Max–Min) labelled counted cycles as "N" rather than "n". This has been corrected throughout the figure.

Proposed manuscript revision:

- Orange box title:
  "Start–Peak–Valley–End stress sequences per window" → "Maximum–Minimum values per window"
- New caption:
  "Figure 3. Extrema sequence methods for reduced-information fatigue assessment. (a) Windowed extraction of start, peaks, valleys, and end (SPVE) points within each window. (b) Concatenation of the windowed SPVE values to form a continuous start–peaks–valleys–end sequence; Rainflow Cycle Counting (RFC) is applied to identify fatigue cycles, including those spanning multiple windows. (c) Resulting fatigue spectrum obtained from the SPVE sequence. (d) Windowed extraction of maximum and minimum stress values per window. (e) Concatenation of windowed maxima and minima to form a maximum–minimum (Max–Min) sequence; Rainflow Cycle Counting (RFC) is applied to identify fatigue cycles. (f) Resulting fatigue spectrum obtained from the Max–Min sequence. The SPVE approach preserves local turning-point structure, while the Max–Min approach provides a highly compressed representation emphasising large-amplitude, long-period stress variations."

- (c) and (f):
  Axis "N" → "n"

Comment 14

Table 1, N_raw, can you use another symbol. N is for number of stresses to failure.

We agree, this is a direct consequence of the notational fix already applied in response to earlier comments (N reserved for cycles-to-failure, n for counted cycles), and N_{raw}/N_{method} created a third, conflicting use of the same letter. We have renamed these to P_{raw} and P_{method} (number of points), which avoids any overlap with the cycle-counting notation.

Proposed manuscript revision:

- Equation (6):
  Before: $Data\ reduction\ [\%] = \left(\frac{N_{raw}-N_{method}}{N_{raw}}\right) \times 100$
  After: $Data\ reduction\ [\%] = \left(\frac{P_{raw}-P_{method}}{P_{raw}}\right) \times 100$
- Definitions (below Table 1):
  Before: N_{raw} is the number of points that belong to the raw signal. $N_{method}$is the number of the points that belong to the sequence signal for each method, with $method = \{L_{f+h}, S, PVEMax - Min\}$.
  After: P_{raw} is the number of points that belong to the raw signal. $P_{method}$is the number of points that belong to the sequence signal for each method, with $method = \{L_{f+h}, S, PVEMax - Min\}$.

Comments 15 and 16

Line 247: $i \in \{low, high\}$ should be in line 246.

Line 247: "relative to the total damage $DLFFD{,}total$" should not be here. Add the symbol D LFFD,total to the end of line 246.

We agree with both points, which together call for the same reorganisation: the domain of the index $i$and the definition of $D_{LFFD,total}$both belong with the bullet defining $D_{i,total}$(line 246), since that is the symbol in which $i$first appears, rather than with the bullet defining $D_{LFFD,i}$(line 247). We have moved both accordingly.

Proposed manuscript revision:

Before:
$D_{i,total}$ is the percentage contribution of component $i$to the total fatigue damage.
$D_{LFFD,i}$, with $i \in \{l, owhigh\}$, is the fatigue damage of component $i$relative to the total damage $D_{LFFD,total}$obtained by the LFFD methodology.
After:
$D_{i,total}$, with $i \in \{low, high\}$, is the percentage contribution of component $i$to the total fatigue damage.
$D_{LFFD,total}$is the total damage obtained by the LFFD methodology.
$D_{LFFD,i}$ is the fatigue damage of component i.

Comment 17

Line 272: If the yaw parameter was used for the conversion to FA/SS, the yaw mean was from a window of 15s and 10min SCADA for the 2 windowing of 15s and 10min, respectively?

The yaw signal used in the FA/SS decomposition is not a SCADA statistical descriptor (i.e. not a 10-minute or 15-second window mean); it is obtained from the high-frequency strain acquisition system at 50 Hz and decimated identically to the strain signals themselves (Section 3). It therefore always matches the native temporal resolution of the strain data being analysed, rather than being drawn from a separately windowed SCADA yaw average, and there is consequently no distinction between the 15-second and 10-minute cases in terms of the yaw signal used. We have clarified this in the text.

Proposed manuscript revision:

- Extending the sentence at Line 272 on FA/SS decomposition:
  "(...) decomposition into fore–aft (FA) and side-to-side (SS) components using the high-frequency yaw signal, decimated identically to the strain measurements and therefore matching the native temporal resolution of the analysed data, and IEC-based wind-speed correction, (…)"

Comment 18

Line 280: Can you define Run, Ready, STOP states shortly?

Briefly: Run indicates the turbine is actively producing power; Ready indicates it is available and grid-connected but awaiting sufficient wind or dispatch conditions to start producing; and STOP indicates the turbine has been intentionally stopped, whether by manual command, scheduled maintenance, or a triggered safety/fault condition. We have added these short definitions at first mention in the manuscript.

Proposed manuscript revision:

- Section 3, extending the sentence at Line 280:
  "The data points are coloured according to the turbine operational status, including Run (actively producing power), Ready (available and connected, awaiting sufficient wind or dispatch conditions to start producing), STOP (intentionally stopped, whether by manual command, scheduled maintenance, or a triggered safety/fault condition), Ambient Cable Unwind, and other less frequent states."

Comment 19

Figure 6: EC3 to be removed from the title of the subplots as they show both EC3 and DNV results.

We have removed "EC3" from all affected panel titles.

Proposed manuscript revision:

- Panel (b): "SD EC3 (log-x, log-y)" → "SD (log-x, log-y)"
- Panel (c): "SD (linear-x, log-y) - EC3" → "SD (linear-x, log-y)"
- Panel (e): "SD EC3 (log-x, linear-y)" → "SD (log-x, linear-y)"
- Panel (f): "SD (linear-x, linear-y) - EC3" → "SD (linear-x, linear-y)"

Comment 20

Figure 7: There is no high 15-sec bar for (b)

The high-frequency bar for panel (b) was indeed missing due to a plotting error. It has been corrected and now correctly shows 100%.

Proposed manuscript revision:

- Figure 7, panel (b), high category: bar corrected from missing/blank to 100%.

Comment 21

Figure 9: Why the SS direction is more sensitive to lower sampling frequencies than the FA direction? What can be the physical explanation?

The likely physical explanation relates to the different loading mechanisms governing each direction. The Fore-Aft (FA) response is dominated by aerodynamic thrust, which varies with wind speed and produces larger-amplitude stress excursions driven by comparatively slow, longer-duration wind-speed variations. This is consistent with the substantially larger total FA damage observed (T1: D_{EC3}= 2.19×10^{-3} (Figure 10) vs. 9.89×10^{-5} for SS (Figure 1 of the current document)) and with the FA stress-range spectrum extending to considerably higher $\Delta\sigma$ values than SS (Figure 10). Because these longer-duration excursions remain adequately resolved even after decimation, FA is comparatively less sensitive to sampling-rate reduction.

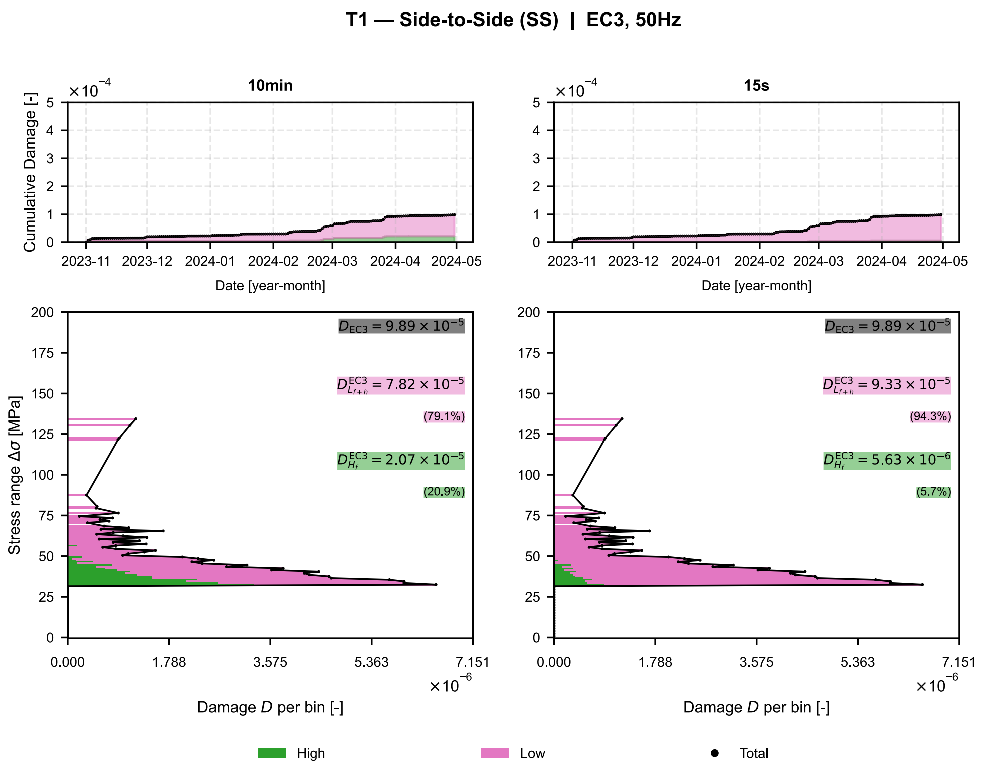

Figure 1 SS-direction fatigue at tower section T1 (six months, 50 Hz), computed using the EC3 detail category 80 σ-N curve (bi-linear formulation, slopes m=3/5). (a–b) Cumulative damage for 10-minute and 15-second windows. (c–d) Final Δσ-D spectra. Total damage is decomposed into within-window full-cycle damage (green) and residual-based damage (pink). The percentages of damage are computed by Equation (2).

The Side-to-Side (SS) response, in contrast, is dominated by rotor torque and rotor-rotational loading, tied to blade-passing frequency and rotor rotation rate rather than to wind-speed variation. This content involves smaller-amplitude, shorter-duration cycles that are more easily lost or distorted when the signal is decimated. Consistent with this, at a decimated sampling rate of 6.25 Hz, the SS direction loses more damage than FA in both the high-frequency (intra-window) and low-frequency (residual) components, for both window lengths: at 10-minute windows, high-frequency loss is −41.1% ± 1.0% (SS) vs. −19.1% ± 0.4% (FA), and low-frequency loss is −17.1% ± 1.5% (SS) vs. −6.9% ± 0.6% (FA); at 15-second windows, high-frequency loss is −40.3% ± 0.9% (SS) vs. −29.0% ± 10.0% (FA), and low-frequency loss is −20.8% ± 1.8% (SS) vs. −9.1% ± 0.8% (FA). The consistently larger low-frequency loss in SS across both window lengths indicates that this shorter-duration, rotor-rotational content is present throughout the SS signal, including within the reconstructed residual sequence, rather than being confined to intra-window cycles alone.

This directional sensitivity is not a fixed property of FA or SS in general: at tower section T2, the pattern reverses, with FA showing greater sensitivity, consistent with that section's proximity to the nacelle, where drivetrain- and rotor-related mechanical excitations introduce comparable shorter-duration content into the FA direction. This confirms that sensitivity to sampling rate follows the local loading mechanism at each section/direction combination, rather than being an inherent property of the FA/SS axis itself.

Proposed manuscript revision:

- Section 4.3, extending the paragraph discussing Figure 9:
  "...indicating higher sensitivity of the SS response to undersampling. This is consistent with the different loading mechanisms governing each direction: FA is dominated by aerodynamic thrust, which varies with wind speed and produces larger-amplitude stress excursions driven by comparatively slow, longer-duration variations, whereas SS is dominated by rotor torque and rotational loading linked to blade-passing frequency, producing smaller-amplitude, shorter-duration cycles that are more susceptible to loss under decimation. This is reflected in the substantially larger total FA damage and its wider stress-range

spectrum (Figure 10), and in the SS direction's greater sensitivity to decimation in both the high- and low-frequency components at 6.25 Hz, for both window lengths (10-minute: −17.1% vs. −6.9%; 15-second: −20.8% vs. −9.1%, for the low-frequency component alone), indicating that shorter-duration, rotor-rotational content is present throughout the SS signal rather than confined to intra-window cycles. This directional sensitivity is not universal: at tower section T2, with FA showing greater sensitivity, consistent with that section's proximity to the nacelle and its associated drivetrain/rotor excitations."

Comment 22

Table 3: Are the storage percentage compared to first row? If yes, first row should be 100 %? And second row should be 107% in total? The mentioned 1 histograms have negligible storage need?

All storage percentages in Table 3 are relative to the raw 50 Hz signal length (consistent with the "Data reduction" definition in Table 1, Equation 6), not to the first row, as the reviewer's reading suggests. We agree this was not stated explicitly, which invited the row-relative interpretation described in the comment. We have added a footnote to the table clarifying the basis of comparison and the negligible cost of a single histogram, without otherwise altering the table.

Proposed manuscript revision:

- Table 3, extending the table caption:
  "All percentages are expressed relative to the raw 50 Hz signal length (Equation 6). A single histogram (~40 stress-range bins) is negligible in storage terms relative to the raw signal length; window-based approaches (rows 1–2) instead require one histogram per 10-minute window, giving ~26,209 histograms over the six-month record."