Discussion of:

Kriging meta-models for damage equivalent load assessment of idling offshore wind turbines

06 September 2025

The manuscript presents a methodology for building surrogate models (metamodels) for predicting loads under idling conditions, using environmental conditions as inputs. The manuscript is well written and easy to follow. Below, some suggestions for improvement:

General comments

1) Novelty: I suggest some more emphasis on the concrete novel elements. Surrogate models are known, including Kriging, and many publications report a few thousand data points as sufficient for model training – and that one shouldn't necessarily use much more as the computational time increases significantly (especially for the Kriging model used by the authors). The novelty must be in studying how to approach the idling problem and if it has differences with the standard "normal operation" modelling approach.
2) I think the authors are mixing together two problems that affect convergence of the time series: the need for the numerical simulator to run-in, and the fact that the time series apparently have non-ergodic properties (impossible to properly characterize statistically) due to the very slow rate of change of some signals. In my view, the significant effort of the authors (tens of thousands of simulations of 60min duration) has in fact mostly been spent on dealing with the non-ergodic properties of the signal, which is a valid problem but unrelated to the run-in time. In fact, the total simulation time might be another factor to study. I suggest the authors to go deeper into this problem as this is a unique challenge for idling conditions.
3) I wonder what is the total contribution from idling events to the lifetime fatigue damage accumulated on a typical site? This could give a very tangible insight into the importance of doing this type of study, and I believe it should be simple for the authors to estimate the number.

Specific comments

4) Abstract: I think the first few sentences in the abstract are unclear. I wouldn't call the surrogate model an alternative to aeroelastic simulations, since one needs to run aeroelastic simulations to train the model. The surrogate model is literally a surrogate or a substitution of the aeroelastic simulations when repeated runs are needed – in that sense, the surrogate models help us by making sure we only need to run the high-fidelity simulation experiment once.
5) Page 2, line 39: IEC 61400-1, ed4 is the current standard. It also considers idling, so I would cite the most recent standard.
6) Page 2, lines 48-50: I think the use of "load cases" here is confusing. I believe the authors mean that approximately 5% of the operational lifetime (5% of the reference 10-minute periods) are spent

under idling conditions. In IEC61400-1 lingo however, “load cases” are specific scenarios (Design Load Cases, DLCs) to be simulated, and they don’t necessarily correspond to a specific time fraction. The “idling” DLC is one – DLC6.4. I suggest replacing “load cases” with “10-minute realizations” or similar.

7) Page 2, lines 54-55: while 5-15% of the lifetime indeed corresponds to more than 100,000 ten-minute periods, I wouldn’t agree that we need 100,000 simulations to cover it as normally the load cases will create a lookup table for a few combinations of wind conditions and apply probability weighting. For example, the normal production load case DLC1.2 normally requires anything between 200 and 1000 simulations, depending on the turbine technology and onshore/offshore configuration. I don’t see why we would need more for idling.
8) Page 4, line 121-123: 10,000 simulations per wind speed to determine run-in times sounds excessive. With 16 wind speed bins, does this mean 160,000 simulations in total? And I believe the 2m/s step is not ideal if we are interested in the idling range, because we have most of the idling happening at very low wind speeds, and we may want higher resolution at those ranges.
9) Section 3.2, line 238: another 10,000 simulations are used to train the Kriging model. Maybe the simulations done for determining the run-in time can be reused instead?