Brief communication: An elliptical parameterisation of the wind direction rose (wes-2024-187)

Response to Reviewer 2

Dear Reviewer,

Thank you for taking the time to consider this submission to WES. I am grateful for your insights, comments and suggestions, and I believe the manuscript will be much improved as a result. Please find my detailed responses below, where I also include your comments in **blue**:

The manuscript proposes a parametric model for the probability of observed wind directions, and it might be used for any circular probability distribution. The aim is to provide a smooth wind direction rose, suitable for optimization of wind farm layouts or advanced wind farm control.

I would disagree with this characterisation of my aim in developing this parametric wind direction rose model. I would instead state the aim as being *the provision of a simple parametric direction rose model which supports standardisation and the identification of generally representative cases, and enables systematic sensitivity analyses of wind rose "shape" impacts on wind farm innovations*. I have, therefore, set out to develop a model which effectively captures the general shape of typical wind direction roses, in much the same way a 2-parameter Weibull distribution captures the general shape of annual mean wind speed distributions. Crucially, neither the Weibull distribution nor the presented parametric wind rose model have much utility if one is working to optimise layout or control at a single real site. In such cases one should simply use the empirical distributions of wind speed and wind direction for the known site. Rather, I believe the presented model has important utility in cases where one is focussed on developing a capability or technique which might then be applied to a variety of sites, and/or if one is seeking to investigate fundamental relationships between site characteristics (including wind direction rose shape) and reliability or yield impacts. An example of the former would be in the development of optimisation tools for layout and control co-design (in which wake effects, and so the wind direction rose, play an important role). In order to demonstrate efficacy for such tools, and motivate their ongoing development or real world application, the potential benefits will generally be quantified for a theoretical wind farm, typically using only a single arbitrarily selected wind rose. If, for example, 1% more power is shown to be generated as a result, that provides some quantification of the potential benefits. However, it is also unclear how much that value might change between sites with more uniform wind roses, versus strongly bi- or uni-directional wind roses. Utilising the presented parametric model, a detailed sensitivity analysis of the wind rose shape impacts on yield may now be undertaken to provide both an improved characterisation of potential benefits (e.g. yield might in-fact vary between 0.8% and 3.4% based on the shape of wind direction rose) and an enhanced conceptual understanding of the problem (e.g. co-design benefits may not be worth pursuing for more uniform wind roses with elliptical parameter below a given value). These same benefits would also hold in analyses of wake impacts on turbine subcomponent reliability, where the goal is not that of characterising reliability impacts for a single real site, but instead to provide a general and fundamental improvement in our scientific understanding of these effects across sites of different types. Beyond this, the parametric direction rose model provides an opportunity to standardise our characterisation of wind rose "shape" and identify normal parameter ranges across which sensitivity analyses should be considered, again in much the same way that the Weibull distribution is used to characterise wind speed distributions.

The basic model takes the shape of an ellipse, and, to allow more flexibility, it folds part of the probability mass in half of the ellipse upon the opposite half. An expression for the area of an ellipse sector is presented and used to fit the parametric model to observed wind sector frequencies. The model-fitting principle is the minimization of the sum of squared errors. For this purpose, the author presents equations for derivatives of the objective function with respect to model parameters. The model does not yet include a directional variation of the wind speed distribution.

I would counter the claim that "The model does not yet include a directional variation of the wind speed distribution". Model fitting does include a determination of the prevailing wind direction, based on the data being fitted to. You are indeed correct that this aspect of modelling fitting is not part of the optimisation scheme. The prevailing wind direction is instead identified by fitting parametric models to the data assuming a prevailing wind direction of a) the circular mean direction and b) the highest probability wind direction bin, before keeping the one which provides the best fit. This heuristic is based on the logic that if neither of these prevailing wind directions produce a good fit to the measured data, then the model is unlikely to result is a good-fit for others. That logic has borne out through testing. While the prevailing wind direction could become an additional optimisation parameter, it would likely introduce problems related to local minima into the optimisation. As the current formulation performs well, and the heuristic has stood up well, I see not immediate need to embed prevailing wind direction within the optimisation itself. Having said that, the above points should probably be more clearly discussed within the paper, and so I will elaborate on these points when revising the manuscript.

Gradient-based layout optimization algorithms will accept larger wake effects in sectors with low frequency of occurrence and thereby smaller contributions to annual energy production. If the input wind rose is too detailed, the algorithm's convergence may be slow, and the solution will be sensitive to random variations. Thus, models with smoother directional variation are needed for some purposes. On the other hand, the wind-rose simplification should not significantly alter the predicted energy production with or without wake correction. At most sites, the wind speed distribution depends on direction, so we risk that the energy production estimate changes after modifications of the wind rose.

These are all valid and important points. As detailed in my first (long) response comment, I am not proposing that this model be applied in cases where design is being undertaken for a single known site. In such cases the empirical distribution should simply be used directly. But you also highlight here a potential application for the parametric model that I'd not previously considered, that of possibly providing a smoother representation of a site's wind rose in order to facilitate a faster first-pass optimisation. That result could then provide a first guess for initialising a second optimisation in which the actual (non-smooth) wind rose is reintroduced. There would of course be caveats to this, such as those you outline. I will seek to include a discussion of both the potential opportunity and the caveats when revising the manuscript. I will also highlight the ever important point that, much like for a Weibull distribution, there will be instances in which more detail is required and so a simplified representation is not suitable.

An ellipse is symmetrical over both major and minor axes, so we might fold over either or both of them. Just remember that the rotation angle should be included as an optimization parameter if we choose to fold over both axes. Unfortunately, the fold-over procedure introduces discontinuities in the dictations along the minor axis, which might reintroduce the disadvantages of the raw wind rose.

I agree you could fold over both axes if looking to extend the model to be more flexible. Having said that, the necessity of increasing the number of model and optimisation parameters by 2 (additional fold + prevailing wind direction) makes we feel the benefits of a simple parametric model might start to be lost, in addition to furthering complicating the optimisation with local minima. But, it is certainly a valid point and I will make sure to include this observation in the revised manuscript. On the point concerning discontinuities, you are right this might slow layout optimisations if smoothness is a primary goal of the parametric model. However, as outlined earlier in my response, smoothness was not the primary motivating factor behind the proposed model.

The model is fitted by a raw wind rose with discrete sector statistics, but it might be more accurate to fit directly by data. The result seems to be a new sector-based distribution, but working with the underlying continuous distribution in optimization algorithms might be better.

The proposed model was developed with a focus on simplicity and easy applicability, which is why it fits directly to the wind rose, rather than raw wind data. Importantly, I believe this still allows all of the principal *aims* (detailed in my first long response comment) of the model to be fulfilled. In addition, there is the added benefit that, in practice, site wind roses are most commonly available in the form of sector probabilities rather than raw wind data.

The von Mises (vM) distribution is the classic model for circular statistics. Due to its unimodal distribution, it is rarely used in wind engineering, but the generalized von Mises distribution (GvM) supports an arbitrary number of modes. GvM models are not easy to fit to data, but Kim and SenGupta present a promising numerical method based on the maximum likelihood principle, see https://doi.org/10.1080/02664763.2020.1796938. The book "Directional Statistics" by Mardia and Jupp discusses more options, see https://onlinelibrary.wiley.com/doi/book/10.1002/9780470316979.

Thank you for highlighting these resources. While there are other approaches to characterising directional distributions, as you correctly point out, they tend to be tricky to fit and can have large numbers of parameters. Additionally, the fitting itself can require a strong knowledge of statistical theory to grasp. Instead of going down that route, I have developed a geometrically driven model which, with a small number of parameters and straightforward sum-of-squares-error fitting, is shown to successfully capture the general shape of various real wind roses. As such, I believe there is significant benefit to the parametric model is have presented.

I once used a more straightforward approach, fitting Fourier splines to the observed sector frequencies and directional variations of the mean and cube of the wind speed. A low-pass filter in wave number domain provided flexible directional smoothing, and Weibull distributions for wind speed from different directions were derived by statistical moments of the wind speed.

This does sound interesting as an approach, but also not very generalizable (i.e. there isn’t a small number of parameters which represents any individual wind rose, such that you can say “most wind roses have the following parameter ranges” etc – which I see as a key benefit of my proposed model). I’d suspect you were aiming to smooth the wind direction rose for improved performance in layout optimisation (or similar)? As described above, while that’s certainly a valuable capability, it was not my particular aim when developing this model.

I suspect that the new elliptical model offers too little flexibility for accurate wind farm production estimates. However, it might be useful for special purposes like the development of wind farm control strategies or fast approximate layout optimization.

As previously described, accurate production estimates for an individual wind farm would indeed not benefit from the proposed model, and instead the empirical distribution should simply be used instead. As you then observe, the parametric model is instead mostly directly conceived as a valuable tool when developing (and exploring potential benefits of) a given capability or technique which might then be applied to a variety of sites, and/or if one is seeking to investigate more fundamental relationships between site characteristics and performance or subcomponent reliability impacts.

Finally, concerning flexibility, more complex wind rose representations can be readily obtained by extending the proposed parametric model to a mixture of such models, allowing for multi-modal wind roses. A drawback of this would be increased numbers of parameters during optimisation, but one could go down a middle-road by requiring the prevailing wind direction of each mixture component to be manually specified. Anyway, within the revised manuscript I will discuss these possibilities and provide an example of a two-model mixture wind rose to highlight that further development towards more general cases is very possible.

SPECIFIC COMMENTS:

P1, line 15: I was puzzled by the expression “energy uplift obtained for a single candidate wind rose”. Try to reformulate for clarity.

Will do!

P2, line 36-48: The explanation of the eccentricity is not used in model formulation, so it might be left out.

I included it for completeness, but will reconsider whether to include or not.

Section 2.3: The multi-case equations in this paragraph are complex to read. Maybe you could simplify by using the Arg or Atan2 functions.

I agonised over the best way to present these equations when developing the paper and this I found to be best. I think there will be multiple cases whichever way the formulation occurs, because of direction bins straddling quadrants. In addition, the presented formulation matches the code implementation (which will be made available alongside the final published paper), and so for consistency reasons I think the current presentation is likely best. Having said that I will take another look to make sure!

Thanks again,

Edward Hart

Senior Lecturer // Chancellor’s Fellow
Wind Energy and Control Centre
Dept. of Electronic and Electrical Engineering
The University of Strathclyde
Glasgow, UK