We have addressed all of the points raised by the reviewer (copied here and shown in red text), and include our responses to each point below (in black text).

Reviewer 1

This is an excellently written paper that has a clear purpose, structure and message. The figures are clear, the method nicely builds on, and quotes, previous work and the conclusions are traceable and of interest. Really interesting to see the result that change from a - to + bias does not follow the diagonal. The discussion on what the non-zero bias at (0,0) and (1,1) implies is a nice way of getting to two key results: that there is a not enough solar radiation reaching the surface when the model correctly predicts clear sky and that there is too much when the model correctly predicts overcast conditions. This is a useful technique for identifying issues with, probably aerosols, and in-cloud water paths. Also interesting to see the result that at this location "clouds are forecast less skilfully in summer, which is when the solar resource is greatest. This paper could probably be accepted as it is, but for thoroughness I include a list of typographical issues and two minor science questions:"
We thank the reviewer for their comments on our submitted manuscript "Evaluating solar radiation forecast uncertainty". Based on the comments and suggestions by the reviewer, we have revised our manuscript.

Minor Issues:
p12, l 33, "a persistence forecast uses the forecast from the day before". Are you sure you don't mean "a persistence forecast uses the OBSERVATIONS from the day before", also I guess these are the "HOURLY observations".
We use hourly forecast values as our persistence forecast. Therefore, we just simply keep the same forecast as was forecast to the previous day to represent the "persistent condition". We updated the sentence to state: "a persistence forecast uses the hourly forecast values from the day before."

p13, l 28 Could you include the formula for the regression that allows you to de-bias your data? I realise that this may only really be applicable at this location and if it were included others may be tempted to apply it elsewhere, so I understand if you would rather not.
As pointed out by the reviewer, our aim in this study was to develop methods applicable to any site. At this particular site, we noted that the relative bias appeared to be constant across a wide range of GHI values, hence the possibility of bias correction. However, this constant bias may not be true at other locations, and would require

further analysis. We decided not to include this bias correction in the manuscript.

Typography:
p2, l 10: suggest changing "by the ECMWF" to "of the ECMWF"
Changed.

p3, l 17: delete comma after "therefore". And remove THE in "do not use these the values".
Changed.

p3, l 25: Kotthaus ref place the (after the name.
Corrected.

p3, l 33: no need for "clearly"
Removed.

p8, lines 11-13 and lines 15-17, these sentences seem like a contradiction (one says you are using a sum (i.e. maximum overlap), then you say random overlap... Do lines 15-17 need to be included at all?
Lines 6–14 refer to the observations, where we explain how the observed cloud cover values are treated. Lines 15-17 refer to the model forecasts. Observations of lcc and mcc are summed due to the nature of observations (time series of cloud hits), whereas we need to combine forecast lcc and mcc.