

METHOD FOR SOLAR IRRADIANCE FORECASTING

TECHNICAL FIELD

[0001] This invention relates to a solar irradiance forecasting method. More particularly, this invention relates to a solar irradiance forecasting method for using a local numerical weather prediction (NWP) model with input from a global NWP model.

BACKGROUND

[0002] The following discussion of the background to the invention is intended to facilitate an understanding of the present invention only. It should be appreciated that the discussion is not an acknowledgement or admission that any of the material referred to was published, known or part of the common general knowledge of the person skilled in the art in any jurisdiction as at the priority date of the invention.

[0003] Solar continues to be the fastest growing renewable technology in the world with an annual average growth rate of over 40% from 2000-2018 for the Organisation for Economic Co-operation and Development (OECD). Many countries across the world now have non-trivial amounts of variable supply in their electricity mix. Solar and wind energy have been added to the electricity grid. The grid operators require solar and wind farm operators to forecast their output for the coming hours or days. Various strategies have been employed to deal with the unknown output from such sources of electricity as solar. In some countries, penalties are imposed for operators who cannot meet the accuracy requirements. In other countries, the grid operators undertake the task of forecasting the output from solar and wind in order to stabilise the grid. Under both these scenarios, the impact from an incorrect forecast can be significant—either financially or in terms of grid stability. It is therefore important to be able to accurately forecast solar irradiance that is closely associated with solar energy output.

[0004] A variety of techniques are commonly used to forecast solar irradiance. These techniques include statistical models, physical models, and combinations thereof depending on the location and the forecast horizon. At time horizons beyond a few hours, techniques that incorporate Numerical Weather Prediction (NWP) output are generally accepted as best practice. During times of cloudiness which prevents the sun's rays from hitting the ground, or at longer time horizons, NWP models can out-perform statistical techniques due to their ability to simulate the atmospheric processes that form

clouds. NWP models operate by discretising the atmosphere across a 3-dimensional grid and stepping forward the governing equations using polynomial expansion. Starting from an initial guess, NWP models can give a rather complete picture of the state of the atmosphere at some later time.

[0005] The Weather Research and Forecasting (WRF) model is a commonly used NWP model for local NWP forecasting. The WRF model is a community-based model that undergoes a major update every year. Each major release incorporates new advancements in the field of NWP, which in recent years has also included physics options/schemes specific to solar forecasting. The use of NWP models, like WRF, for forecasting irradiance at the local scale is gaining attention. Broadly speaking the WRF model, like all NWP models, suffers from error spreading due to three main sources: sub-grid scale processes not captured by the resolution of the WRF model, complex processes that require too much computation to be resolved explicitly, and errors in the initial state of the WRF model. The first two error sources can be partly addressed through physics model parameterisations—like the Radiative Transfer Model (RTM) or microphysics schemes. Parameterisations attempt to mimic the behaviour of a specific process and can feed the result of this behaviour back to the WRF model. One approach taken for reducing model error in WRF is to choose a set of parameterisations that best resolves the processes associated with a region. It is, however, observed that using such a single best configuration, biases or errors still exist between forecasted and observed solar irradiance.

[0006] There is therefore a need for a system which addresses, at least in part, one or more of the foregoing problems.

SUMMARY

[0007] According to an aspect of the present disclosure, there is provided a method of forecasting solar irradiance for a selected day using a local NWP model based on forecast output of a global NWP model. The method includes adjusting moisture values in the forecast output of the global NWP model by an adjustment amount corresponding to a window period including the selected day, and forecasting solar irradiance is based on the forecast output with moisture values adjusted.

[0008] In some embodiments of the method, the window period is a calendar month.

[0009] In some embodiments of the method, adjusting moisture values includes adjusting water vapour mixing ratios.

[0010] In some embodiments of the method, adjusting water vapour mixing ratios includes converting each water vapour mixing ratio into a corresponding relative humidity value, adjusting the relative humidity value by the adjustment amount, and converting the adjusted relative humidity value back to a water vapour mixing ratio.

[0011] In some embodiments of the method, the method further includes nudging forecasted values of the local NWP model towards adjusted or non-adjusted global NWP model values.

[0012] In some embodiments of the method, the method further includes determining an adjustment amount for each window period of a plurality of window periods.

[0013] In some embodiments of the method, each adjustment amount is determined empirically.

[0014] In some embodiments of the method, determining an adjustment amount for each window period only when there is a systematic bias in previously forecasted solar irradiance for the window period.

[0015] In some embodiments of the method, the method further includes writing the adjusted moisture values to a new intermediate file for use by a metgrid program.

[0016] According to another aspect of the present disclosure, there is provided a computer device for forecasting solar irradiance for a selected day using a local NWP model. The computer device a processor that is operable to adjust moisture values in the forecast output of the global NWP model by an adjustment amount corresponding to a window period including the selected day, and forecast solar irradiance is based on the forecast output with moisture values adjusted.

[0017] In some embodiments of the computer device, the window period is a calendar month.

[0018] In some embodiments of the computer device, moisture values are adjusted by adjusting water vapour mixing ratios.

[0019] In some embodiments of the computer device, wherein the processor adjusts water vapour mixing ratios by converting each water vapour mixing ratio into a corresponding relative humidity value, adjusting the relative humidity value by the adjustment amount, and converting the adjusted relative humidity value back to a water vapour mixing ratio.

[0020] In some embodiments of the computer device, the processor further nudges forecasted values of the local NWP model towards one of adjusted and non-adjusted global NWP model values.

[0021] In some embodiments of the computer device, the processor further writes adjusted moisture values to a new intermediate file for use by a metgrid program.

[0022] Other aspects and advantages of the invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

BRIEF DESCRIPTION OF DRAWINGS

[0023] The invention will be better understood with reference to the drawings, in which:

Figure 1 is a flow diagram showing a method of forecasting solar irradiance according to an embodiment of the invention;

Figure 2 is a flow diagram showing a method of obtaining day-type specific configurations of physics schemes according to an embodiment of the invention for use with the forecasting method in Figure 1;

Figure 3 is table of daily average K_t binning, wherein a maximum value defines an upper limit of a bin, the date chosen with corresponding K_t value for each bin to be used in the method of Figure 2. The dates are in local time in Qinghai;

Figure 4 is a graph showing observed solar irradiance for the dates in Figure 3 in time order;

Figure 5 is a table of WRF surface layer schemes that is considered in the method in Figure 2;

Figure 6 is a table of land-surface model schemes that is considered in the method in Figure 2;

Figure 7 is a table of PBL schemes that is considered in the method in Figure 2;

Figure 8 is a histogram of RMSE performance for all 15 days and for all 179 surface and PBL scheme combinations that is obtained using the method in Figure 2;

Figure 9 is a table of microphysics options that is considered in the method in Figure 2;

Figure 10 is a table of cumulus parameterization options that is considered in the method in Figure 2;

Figure 11 is a table of RTM options that is considered in the method in Figure 2;

Figure 12 is a histogram of RMSE performance for the RTM, microphysics and cumulus optimization using all 15 test period days that is obtained using the method in Figure 2;

Figure 13 is a table showing a best overall WRF configuration the 15-day test period and optimum configurations for different day types;

Figure 14A is histogram of results for the RTM, microphysics and cumulus combinations for sunny days;

Figure 14B is a graph showing solar irradiance for the sunny days obtained using the optimum sunny day configuration and best overall configuration, together with observed solar irradiance;

Figure 15A is histogram of results for the RTM, microphysics and cumulus combinations for partly-cloudy days;

Figure 15B is a graph showing solar irradiance for the sunny days obtained using the optimum partly-cloudy day configuration and best overall configuration, together with observed solar irradiance;

Figure 16A is histogram of results for the RTM, microphysics and cumulus combinations for cloudy days;

Figure 16B is a graph showing solar irradiance for the sunny days obtained using the optimum cloudy day configuration and best overall configuration, together with observed solar irradiance;

Figure 17 is a graph showing solar irradiance for all 15 days obtained using the best overall configuration, a combined WRF (using a combination of optimum configurations for the different day types) and EC model output;

Figure 18 is a table showing RMSE and Bias values for all 15 days for the outputs in Figure 17;

Figure 19 is a scatter plot showing EC and observed daily averaged K_t values for 2018;

Figure 20 is table showing RMSE and Bias values for a 3-month validation set. Results based on observed type of day is also shown;

Figure 21 is a table of solar irradiance for cloudy days of the 3-month validation set in Figure 20;

Figure 22 is a table showing forecasting performance of a WRF model for the Xinjiang region using combined day-type specific configurations, best overall configuration and the EC model;

Figure 23 is a graph showing solar irradiance obtained using the best overall configuration, from the EC model and observed values;

Figure 24 is a histogram of RMSE performance for all configurations/combinations of physics schemes for all 15 test days;

Figure 25 is a table of configurations that resulted in lowest RMSE for all 15 days and for days of each day-type;

Figure 26 is a graph of solar irradiance obtained using the best overall configuration and sunny-day configuration, and observed values;

Figure 27 is a graph similar to Figure 26 but for cloudy days;

Figure 28 is a graph of solar irradiance for all 15 days obtained using the best overall configuration, a combined WRF output and EC model output;

Figure 29 is a flow diagram showing a method of addressing bias in the forecasting in Figure 1;

Figure 30 is a flow diagram showing a method of adjusting moisture values in the forecast output of a global NWP model that is used in the method in Figure 29;

Figure 31 is a scatter plot showing EC and observed daily averaged K_t values for May 2018 to July 2019;

Figure 32 is a graph of monthly averaged bias for cloudy days in a testing;

Figure 33 is a graph of solar irradiance obtained using the optimum cloudy day configuration with 12% RH adjusted input to the WRF model;

Figure 34 is a table showing statistics for January 2019 EC forecasted cloudy days; and

Figure 35 is a block diagram of a computer device according to an embodiment of the invention suitable for performing any one of the above-mentioned methods.

DETAILED DESCRIPTION OF THE EMBODIMENTS

[0024] Throughout this document, unless otherwise indicated to the contrary, the terms “comprising”, “consisting of”, “having” and the like, are to be construed as non-exhaustive, or in other words, as meaning “including, but not limited to.”

[0025] Furthermore, throughout the specification, unless the context requires otherwise, the word “include” or variations such as “includes” or “including” will be understood to imply the inclusion of a stated integer or group of integers but not the exclusion of any other integer or group of integers.

[0026] Unless defined otherwise, all technical and scientific terms used herein have the same meaning as is commonly understood by a skilled person to which the subject matter herein belongs.

[0027] As shown in the drawings for purposes of illustration, the invention may be embodied in a more accurate method of forecasting solar irradiance for a selected day using a local NWP model based on forecast output of a global NWP model. The method includes adjusting moisture values in the forecast output of the global NWP model by an adjustment amount corresponding to a window period including the selected day, and forecasting solar irradiance is based on the forecast output with moisture values adjusted.

[0028] Specifically, Figure 1 shows a method 2 for forecasting solar irradiance of a region on a selected day. Solar irradiance is the power per unit area received from the sun in the form of electromagnetic radiation as measured in the wavelength range of the measuring instrument. The solar irradiance is measured in watt per square metre (W/m^2) in SI units. The method 2 includes forecasting 4 solar irradiance using a local numerical weather prediction (NWP) model, such as but not limited to, a Weather Research and Forecasting (WRF) model. The basic idea of NWP is to sample the state of the atmosphere at a given time and use the equations of fluid dynamics and thermodynamics to estimate the state of the fluid at some time in the future. As known to those skilled in the art, the WRF model is a mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting applications. It features two dynamical cores, a data assimilation system, and a software architecture supporting parallel computation and system extensibility. The WRF model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers. The WRF model can produce simulations based on actual atmospheric conditions (i.e., from observations and analyses) or idealized conditions. The WRF model offers operational forecasting a flexible and computationally efficient platform, while reflecting recent advances in physics, numerics, and data assimilation contributed by developers from the expansive research community.

[0029] The WRF model takes as input the forecast output 6 of a global NWP model, such as but not limited to, the European Centre for Medium-Range Weather Forecasts (ECMWF) model, which is also commonly known as the European model or EC model. The EC model includes both a "deterministic forecast" mode that is a single model run that is relatively high in resolution as well as in computational expense. The EC model's **output 6** provides initial condition (IC) and lateral boundary conditions (LBC) for the WRF model.

[0030] As known to those skilled in the art, the WRF model supports a large variety of physics schemes/options, various lateral conditions, data assimilation one-way, two-way nesting and moving nests. The WRF model has a main physics interface (not shown) that includes:

- microphysics: Schemes ranging from simplified physics suitable for idealized studies to sophisticated mixed-phase physics suitable for process studies and NWP;
- cumulus parameterizations: Adjustment and mass-flux schemes for mesoscale modelling;
- surface physics: Multi-layer land surface models ranging from a simple thermal model to full vegetation and soil moisture models, including snow cover and sea ice;
- planetary boundary layer (PBL) physics: Provides boundary layer fluxes including turbulent energy budget and vertical diffusion in whole column; and
- radiative transfer model (RTM): Longwave and shortwave schemes with multiple spectral bands and simple shortwave scheme suitable for climate and weather applications; cloud effects and surface fluxes are included;
- turbulence/diffusion: Sub-grid effects on all fields including computation of advection coefficients.

[0031] Prior to use of the WRF model to forecast 4 solar irradiance, different configurations of physics schemes are determined for different types of days. For example, days may be classified according to cloud cover into cloudy days, partly-cloudy days and sunny days. There will be an optimum configuration of physics schemes corresponding to each day type although any suitable configuration will do. For cloudy days, there will be a cloudy day configuration of physics scheme. For partly-cloudy days,

there will be a partly-cloudy day configuration of physics scheme. And for sunny days, there will be a sunny day configuration of physics scheme. The different day-type configurations of physics schemes may be the same or different for the different types of days. The different day-type configurations of physics schemes may be the same or different across different locations.

[0032] From the forecast output 6 of the EC-model, the type of day for the selected day is obtained 8. The optimum configuration of physics schemes corresponding to the day type of that day is then selected 10. Forecast 4 of the solar irradiance for that day is then carried out using the WRF model with the selected configuration of physics schemes. It will be shown later that forecast of solar irradiance using a configuration of physics schemes specific to the type of day is more accurate than that obtainable with the use of a generic configuration across all day-types.

[0033] A method 20 for obtaining the above-described day-type specific configurations of physics schemes for a first region, e.g., Qinghai in China, is next described. The Qinghai region is mid-latitude and largely at high elevation. The Qinghai region exhibits seasonal variability that includes surface temperatures well below 0 degrees Celsius in the winter and warm mostly dry conditions in the summer. The method 20 starts in a SELECT IRRADIANCE DATA step 22, wherein observed solar irradiance data of a Typical Meteorological Year for Qinghai is selected. The method 20 next proceeds to a CONVERT IRRADIANCE TO K_t step 24, wherein the irradiance data is converted to sky cover amounts, k_t . To avoid spurious values for k_t , which are common near sunrise and sunset, only irradiance values with a solar zenith less than 70 degrees were considered and converted each day. The sky cover amount, k_t , is calculated by dividing the observed or modelled irradiance by the theoretical clear-sky irradiance as follows:

$$k_t(t) = \frac{I_{observed}(t)}{I_{clearsky}(t)}, \quad (1)$$

where, $I_{observed}$ is the observed or modelled irradiance for the Qinghai site, and $I_{clearsky}$ is the theoretical clear-sky irradiance as calculated by using the default Ineichen/Perez model from the python package pvlib, known to those skilled in the art. The method 20 next proceeds to an OBTAIN DAILY AVERAGE K_t step 26, wherein a daily average k_t is obtained from the k_t values.

[0034] The method 12 next proceeds to a SORT DAYS step 28, wherein the days in the year are then sorted in ascending order of daily average sky cover amount, i.e. average k_t values. The method 20 next proceeds to a SPLIT DISTRIBUTION step 30, wherein the sorted days are divided into x number of bins, e.g., 15 bins, of equal weight. Each bin can include multiple days. The method 20 next proceeds to a CALIBRATION step 32, wherein it is checked if there is a large bias in the distribution. If the bias isn't large, like those obtained for the Qinghai region, the first five bins with days having the lowest daily average sky cover amount are considered to represent cloudy days. The next five bins include days that are considered partly cloudy. And the final five bins with days having the highest daily average sky cover amount are considered to represent sunny days. It should be appreciated that more or less number of day-types may be used although only three day-types are considered here. The method 20 next proceeds to a CHOOSE A DAY FROM EACH BIN step 34, wherein a date is chosen from each of the fifteen bins and, where possible, choices were made to maximise the representation from across the year. The date close to or at the centre of each bin may for example be chosen. Figure 3 shows the fifteen bins of days in 2018 and the fifteen days chosen from the bins. Out of the fifteen days chosen, the first five days with the lowest k_t values are classified as cloudy days, the next five days with intermediate k_t values are classified as partly cloudy days and the last five days with highest k_t values are classified as sunny days. Figure 4 shows the corresponding time series plot of solar irradiance for the fifteen chosen dates in 2018.

[0035] The method 20 next proceeds to a RUN NWP step (not shown), wherein numerical weather prediction (NWP) using the WRF model is run using different configurations of physics schemes for the five days of each day-type. More specifically, the WRF model version 3.9.1 with initial and boundary conditions input from the 1200 UTC dataset provided by the European Centre for Medium Range Weather Forecasting (ECMWF), hereafter referred to as the EC model, is used. Inputs from the EC model are provided at a 3-hourly interval and with spatial resolution of roughly 0.1°. The WRF model is set up within the EC model grid with one domain at 3km spatial resolution and with 65 vertical layers. Solar irradiance is forecasted using the WRF model for 52 hours from each 1200UTC dataset release from the EC model, and utilising the day-ahead hours of +28 to +52 hours, which corresponds to midnight to midnight local time in

Qinghai. The forecast location is in the centre of the WRF domain to reduce the influence of interpolation at the model domain boundary.

[0036] As described above, the WRF model include several physics schemes. The physic schemes that impact cloud formation and cloud type are considered in order to find schemes that can most appropriately represent the cloudy processes at the selected location. The microphysics schemes that estimate cloud particle concentration and cumulus schemes that estimate the movement of these cloud properties through convective processes are likely candidates for optimization. Parameterisation of surface processes may also be considered. Specifically, those processes relating to the exchange of moisture and energy from the surface, i.e. surface physics and land-surface schemes, as well as the upward mixing of the atmosphere due to the roughness of the surface, i.e. PBL schemes, are all considered.

[0037] In the RUN NWP step, the impact of the surface and PBL schemes on the performance of WRF day-ahead forecast of solar irradiance, more specifically, global horizontal irradiance (GHI) is first evaluated. GHI is the total irradiance from the sun on a horizontal surface on earth. It is the sum of direct irradiance (after accounting for the solar zenith angle of the sun and diffuse horizontal irradiance. The surface and PBL schemes utilised are listed in Figures 5-7. A total of 179 combinations are tested (less than the complete set of all possible combinations due mostly to incompatibility of schemes and/or computation constraints). During testing of the surface and PBL schemes, the RTM, microphysics and cumulus schemes are kept constant and set to RRTMG long-wave, Dudhia short-wave, Thompson microphysics and no cumulus parameterisation respectively.

[0038] The method 20 next proceeds to COMPARE RESULTS WITH OBSERVERED VALUES step (not shown), wherein the forecasted results, more specifically the WRF SWDOWN (short-wave global horizontal irradiance), are compared with observation values via two measures: Root Mean Square Error (RMSE) and Bias. These measures are defined as follows:

$$RMSE = \sqrt{\frac{1}{N} \sum (I_{pred} - I_{obs})^2}, \quad (2)$$

$$Bias = \frac{1}{N} \sum (I_{pred} - I_{obs}), \quad (3)$$

where, I_{pred} is the WRF/EC predicted irradiance (GHI) and I_{obs} is the observed/measured GHI at the Qinghai site. The metrics are calculated at the 15-minute frequency of the observations. Irradiance values are determined by the WRF model RTMs at 3-minute frequency but only the 15-minute values starting on the hour are analysed in accordance with the observations. Night-time values are determined to be when the observations are equal to zero and these times are not included in the RMSE and Bias calculations.

[0039] Where appropriate, EC model SSRD (surface solar radiation downwards) values are also compared to the observation values. The EC model irradiance values are 3-hourly accumulated by the WRF model from initialisation and in units of J/m^2 , which requires de-constructing before comparing to the observation values. The average W/m^2 values from each 3-hour window are calculated by dividing by the number of seconds in 3 hours, which is done for both the EC model and 3-hourly accumulated clear-sky model values. The average k_t values are then calculated following Equation (1) and a linear interpolation performed in k_t -space to achieve 15-minute k_t values. Finally, 15-minute EC model GHI values are extracted by multiplying by the 15-minute frequency clear-sky model values.

[0040] After the forecast/predicted results and RMSE are obtained, a histogram of results, based on RMSE, for the 179 combinations of physics schemes for all 15 days is plotted and shown in Figure 8. As can be seen from Figure 8 there is very little dependence on the surface or PBL schemes for the performance of WRF day-ahead forecast of GHI. The difference between the best performing combination and the worst performing combination was about 0.5% based on RMSE. The small cluster centred on about 270.3 W/m^2 involved all except one of the combinations of NCEP Global Forecast System surface layer scheme and either RUC or 5-layer thermal diffusion land surface model. Therefore, for the Qinghai site, any of the tested combinations of surface and PBL schemes can be used without expecting the choice to significantly influence forecasting performance. Nevertheless, when forecasting next using different combinations of RTM, microphysics and cumulus schemes, the optimum combination of NCEP Global Forecast System surface layer scheme, RUC land surface model and UW (Park and Bretherton) PBL scheme is utilised.

[0041] Solar irradiation forecast with different RTM, microphysics and cumulus schemes are carried out in the RUN APP step. The different options for the RTM,

microphysics and cumulus parameterisation are outlined in Figures 9-11. During testing of these schemes, the surface and PBL schemes are kept constant at the optimum as described above. A total of 160 combinations of RTM, microphysics and cumulus schemes are run with 120 successfully simulating all fifteen days and subsequently being analysed. The configurations failing to complete are likely due to incompatibility of scheme options.

[0042] With respect to the RTM schemes, the long-wave version of those listed in Figure 11 are used, with the exception of the Dudhia short-wave, which has no long-wave counterpart and so the RRTMG long-wave scheme is utilised instead. In deciding which cumulus schemes to use, preference is given to the schemes that has some scale-awareness. At 3km horizontal resolution it is generally considered unnecessary to use cumulus physics as convective processes start to become explicitly resolved by the model. Other options that relate to cloud processes and their relationship to radiation that are tested include sub-grid scale cloud interaction ("cu_rad_feedback" in WRF), which is available and was turned on for Grell and Kain-Fritsch cumulus schemes, cloud effects on optical depth ("icloud" in WRF), which is available and turned on for Dudhia and RRTMG-based RTM models, and feedback from shallow cumulus to the RTM ("shallow_cu_forced_ra" in WRF), which is available for Kain-Fritsch based cumulus schemes and was used with the default binning (21 bins) of potential temperature and mixing ratio.

[0043] Figure 12 shows a histogram of RMSE based on all fifteen days and for all combinations of RTM, microphysics and cumulus parameterisation schemes that successfully completed all fifteen simulation days. A total of 120 schemes are included in Figure 12. Compared to the surface and PBL optimisation, the RTM, microphysics and cumulus combinations show much more variability in forecasting performance. It should be noted that the cluster of combinations at around 295 W/m² RMSE all involved the use of the CAM RTM model, which appears inaccurate and therefore inappropriate for the Qinghai site. The top performing combinations from optimisation using all fifteen test period days is outlined in Figure 13.

[0044] The method next proceeds to a CHOOSE DAY-TYPE SPECIFIC CONFIGURATIONS step (not shown), wherein the full set of fifteen days is split into their respective cloudy, partly cloudy and sunny categories and the RMSE analysis is run for days for each day-type category. The result is compared to that with the configuration

that is found to work best overall for all fifteen days. Figure 14A shows a histogram of RMSE performance for the different combinations of RTM, microphysics and cumulus options for the five sunny days. Figure 14B shows the time-series plot comparing the best performing sunny day configuration with the best overall configuration for the five sunny days.

[0045] For the five sunny days, the RMSE using the best sunny configuration is 69.2 W/m² while the RMSE using the best overall configuration is 70.02 W/m². The difference between the two RMSEs is small. The reason for the similarity in RMSE can be traced back to the WRF configurations, which are almost identical as shown in Figure 13. Clustering in performance was evident for the sunny days as seen in Figure 14A. Figure 14A shows that there are three distinct clusters. The difference in configuration between the first cluster at around 80 W/m² and the second cluster beyond 100 W/m² occurred when switching to Thompson microphysics scheme with RRTMG long-wave and short-wave schemes in conjunction with certain cumulus options. The jump in performance from cluster two to cluster three was for similar reasons but with Dudhia short-wave scheme instead of RRTMG.

[0046] For partly cloudy days the difference was more marked, although most of the difference was due to one of the five partly-cloudy days, i.e. 3 Nov 2018. Cumulus physics was preferred for partly cloudy days as shown in Figure 13, resulting in a more accurate representation of the cloud amount on day four (the second cloudiest from the partly cloudy category) (Figure 15B). The RMSE for the best partly cloudy configuration on partly cloudy days was 197.31 W/m² compared to 214.49 W/m² for the best overall configuration on the same days.

[0047] The biggest difference between choosing the best overall WRF configuration versus a configuration chosen based on day-type is for the cloudy category. Cumulus physics are also preferred for the cloudy category (Figure 13) and the difference in RMSE is 15.7% when comparing the best cloudy day configuration (314.29 W/m²) with the best overall configuration (372.9 W/m²). Figure 16B shows that this difference is also much more visible in the time series across nearly all the five cloudy days. Some clustering in performance is also visible for the cloudy category (Figure 16A). The first cluster (at around 360 W/m²) for cloudy days is due to the configurations using RRTMG long-wave and Dudhia short-wave RTMs while the second cluster (at around 420 W/m²) is centred on the configurations using the CAM RTM.

[0048] Advantageously, the forecasting is more accurate when the WRF RTM, microphysics and cumulus schemes are tailored for the type of day. For the 15-day test period when the optimum configuration for each day-type was chosen, not only was the overall RMSE lower than that obtainable with the best overall configuration, the RMSE was also lower than the EC global model input. Figure 17 illustrates this difference for the 15-day time series while Figure 18 lists the RMSE and bias statistics for the 15-day optimisation period. The combined WRF represents a concatenation of optimum sunny, partly-cloudy and cloudy configurations for their respective days.

[0049] The day-type specific configurations obtained are further tested on a three-month validation set (January-March 2019). As an observed day in this three-month period is not known a priori, the forecast from the EC model is used to determine the type of the observed day so that the corresponding optimum WRF physics configuration for that day may be chosen. While this may introduce a source of error as the EC model forecast is not 100% accurate on the type of day, it is a practical choice. Figure 19 shows a scatter plot of EC-derived daily average K_t versus the observed daily averaged K_t for 2018. There is a reasonable correlation at 0.746 between results of the EC model and observed daily averaged cloud amount with an expected sunny bias from the EC model.

[0050] For each day in the validation set the EC-derived categories for sunny, partly cloudy and cloudy are used to determine the physics configuration for the WRF model. If the EC model forecasted a sunny day then the sunny day WRF physics configuration is used, and so on for the other two categories. The best overall WRF configuration is also run in parallel for each day as a comparison. The resulting statistics are outlined in Figure 20.

[0051] For the validation set the use of adaptive physics, i.e. choosing a day-type configuration of physics schemes corresponding to the type of day, is clearly superior than the use of a single WRF configuration for all days. Both the RMSE and bias statistics are favourable for the use of adaptive WRF physics. While the EC model out-performs the combined model for Jan-Mar 2019 this was dependent on the type of day and not true for the cloudy category. For cloudy days in Jan-Mar 2019 the use of adaptive physics remains better than a single WRF configuration, and with smaller error than the EC model. Figure 21 shows the time series of Jan-Mar 2019 cloudy days based on the observed conditions. For cloudy days it is evident that the overall best WRF configuration greatly over-predicts cloud amount with bias of -78 W/m^2 compared with the adaptive

WRF configuration at -6.7 W/m^2 and the EC model at 132.68 W/m^2 , which greatly over-predicts irradiance as shown in Figure 20.

[0052] The RMSE and bias statistics for the validation set split by day type are also outlined in Figure 20. As can be seen from the day type results the adaptive physics is a better choice overall and for any type of day while the EC model performance deteriorates with increasing cloud amount. For cloudy conditions it is generally expected that limited-area physical downscaling produce better resolved clouds and cloud processes through tailored physics and higher resolution for the Qinghai site.

[0053] The above method 20 is repeated for a second region, the Xinjiang region in China, for days in 2018 and 2019 to find optimized WRF model configuration of physics schemes for the different types of days, i.e., sunny days, partly-cloudy days and cloudy days. The observations at 15-minute intervals have been filtered for erroneous data through a system of checks that examine various known errors including non-zero night-time values, shading of the sensor leading to unrealistic daily maximum values, day-time values that are unchanging, day-time values of zero, and days with mostly missing values.

[0054] Since physics options including Planetary Boundary Layer (PBL), surface layer physics and land-surface model is shown to have little or no effect on the day-ahead forecast of irradiance for the Qinghai region, they are not considered for the Xinjiang region. The best overall configuration (the configuration with lowest RMSE for all 15 day regardless of day types, outlined in Figure 22) is plotted in Figure 23, while a histogram of RMSE for the various configurations is shown in Figure 24. From the histogram in Figure 24, it can be seen that there is a cluster of similarly performing near-optimum configurations at approximately 150 W/m^2 RMSE. Beyond this cluster there is a set of configurations greater than 325 W/m^2 , which all involved the CAM RTM model.

[0055] When the analysis of RMSE was split based on the type of day, sunny, partly cloudy and cloudy, the best performing configuration listed in Figure 25 showed some changes compared to those obtained for the Qinghai region, and the overall RMSE was reduced when the configuration was tailored and used in a combined model. The combined model is created by concatenating the day-type configurations (5 days from each day-type). In this way the combined model was a best-case scenario if the type of day were known in advance and the WRF model configuration chosen according to the best performing for that day type.

[0056] Figure 26 outlines the time series of sunny days from the optimisation set of fifteen days. Figure 26 compares the best sunny day configuration with the best overall configuration for the five sunny days. As can be seen from Figure 26, the improvement by using a physics combination tailored to sunny days was marginal and for the five sunny days the difference in RMSE is merely 5.9% (83.8 W/m^2 compared to 79.9 W/m^2).

[0057] For partly cloudy days, the best partly-cloudy day configuration is the same as the best overall configuration (see Figure 25). However, for cloudy days the difference between the tailored and best overall configurations is more obvious as shown in the results in Figure 27. For cloudy days there is a 15% improvement in RMSE when the best cloudy day configuration is selected instead of the best overall configuration (171.5 W/m^2 compared to 145.8 W/m^2 for the cloudy day configuration), with most of the improvement coming on the second day.

[0058] Comparing the configurations for the day types with the best overall (Figure 25) it may be a surprising result that a cumulus scheme was preferred for sunny days and not for cloudy days for the Xinjiang region. However, for sunny days the top 80 configurations are separated by just 3.8 W/m^2 RMSE and not all the configurations utilise cumulus physics, suggesting little dependence on this option for sunny days (not shown). For cloudy days there was a much larger spread in performance. By comparison, the top 80 cloudy configurations are separated by 153 W/m^2 RMSE and the top 10 configurations are separated by 4.4 W/m^2 RMSE with eight of the top 10 utilising some form of cumulus physics (either shallow or not) (not shown). Comparing the combined model with the best overall configuration also saw some improvement, mostly in the cloudy days (Figure 28). RMSE is reduced by 7.4% for the combined model when compared to the best overall model, with the combined model performing approximately as well as the EC model (Figure 22).

[0059] The cloudy days in the optimisation set showed the best performance increase when the physics was tailored for that type of day. The performance of this configuration over a longer time period is explored. Cloudy days from May 2018 to July 2019 (hereafter referred to as the testing set) and where cloudy days were defined using the EC model forecast are analysed. The EC model is used for operational considerations. As mentioned previously, the type of day being forecasted is not known a priori. The correlation between the type of day being forecast by the EC model and the observed day during this period was quite high (0.81 based on daily averaged k_t values) and the

scatter plot of daily averaged k_t shows good relationship as shown in Figure 31. Figure 31 suggests that little error is introduced when using the EC model to forecast the type of day.

[0060] In order to classify the days into cloudy, partly cloudy and sunny a similar process to the physics optimisation set described earlier is used. The distribution of EC daily average k_t values were split into 15 equal weighted bins with increasing k_t value. As is rather evident from Figure 31 the EC model had a sunny bias for the testing set. The number of bins in the cloudy versus other categories is thus adjusted. The first three bins for EC are used to define cloudy with a cut-off daily averaged k_t value of 0.6382. The cloudy day configuration is run for all EC model forecasted cloudy days for the testing set and the monthly averaged bias is shown in Figure 32.

[0061] From Figure 32 it can be seen that forecasted cloudy days, even with the use of the optimized cloudy day configuration, still exhibited very large negative biases during the wintertime (too cloudy). Months other than November through March also had a negative bias tendency but with much smaller magnitudes. A method 50 shown in Figure 29 for addressing the wintertime bias, using January 2019 as an example, is next described.

[0062] The method 50 starts in a COMPARE WITH OBSERVATIONS step 52, wherein forecast results for cloudy days are compared with in-situ observation values. The method 50 next proceeds to an IDENTIFY BIAS step 54, wherein it is determined if forecasting using the optimised configuration of the WRF model still exhibits bias and the bias is systematic. If it is determined that bias exists and the bias is systematic as determined in the IDENTIFY BIAS step 54, the method proceeds to a ADJUST MOISTURE step 56, wherein bias is addressed by way of adjusting the 3D moisture field that is input into the WRF model by an adjustment amount. Moisture in the form vapour, cloud condensation nuclei, or other forms influences the surface shortwave irradiance by virtue of blocking/deflecting or absorbing and re-emitting incoming shortwave radiation from the sun. Under normal circumstances the WRF model is used in such a way that the limited area domain is initialised by the global/parent model and the only extra information at following time steps comes in the form of boundary conditions from the global/parent model. However, functionality exists whereby the WRF model can be nudged towards the global/parent model values (grid analysis nudging, called grid fdda in WRF). The method 50 next proceeds to an NWP SIMULATION step 58, wherein WRF

forecasting is performed with nudging. By using reasonable nudging coefficients, and where necessary by adjusting the moisture values from the global model, it is shown that surface irradiance biases can be effectively reduced.

[0063] The manner in which moisture values are adjusted is next described with the aid of Figure 30. The EC model conditions that is used to drive the WRF model come in a GRIB data format and at 3-hourly intervals. From the GRIB data, a WPS module unpacks the data and creates a series of intermediate files from Fortran write statements. The intermediate files are then ingested by the metgrid program where the data are spatially interpolated onto the WRF model grid. The water vapour mixing ratio values contained in the indeterminate files are overwritten for circumstances where the optimised WRF model is determined to be systematically biased in surface irradiance. The method uses Fortran code to read in the water vapour mixing ratio (Q_{vap}), pressure and temperature values on the EC model grid, calculate Relative Humidity (RH), adjust the RH values and then write out altered water vapour mixing ratios (Q_{vap}) to a new file. RH is utilised instead of directly altering the Q_{vap} due to RH being bounded from 0-100, which leads to a lower susceptibility for extreme alteration. RH is calculated at each grid point and EC model level as follows:

$$RH (\%) = 100.0 \times \frac{P_{vap}}{P_{sat,vap}} \quad (4)$$

where, the vapour pressure, P_{vap} , is

$$P(Pa)_{vap} = \frac{Q_{vap} \times P}{\left(\frac{R_{gas}}{R_{vap}} + Q_{vap}\right)} \quad (5)$$

where R_{gas} is the gas constant for air (287 J/Kg*K), R_{vap} is the specific gas constant of water vapour (461.6 J/Kg*K) and the $P_{sat,vap}$ is the saturation vapour pressure calculated following the data on page 113 of Curry and Webster 1999 as:

$$P(Pa)_{sat,vap} = 100.0 * 6.11 \times e^{\frac{L}{R_{vap} \times \left(\frac{1}{273.15} - \frac{1}{T}\right)}} \in T < 223.15K \quad (6)$$

$$P(Pa)_{sat,vap} = 100.0 * \left[a(1) + \sum_{i=2}^7 a(i) \times (T - 273.15)^{i-1} \right] \in T \geq 223.15K \quad (7)$$

where L is the latent heat of vapourisation ($2.5e^6$ J/K), and $a = (6.1117675, 0.443986062, 0.143053301e^{-1}, 0.265027242e^{-3}, 0.302246994e^{-5}, 0.203886313e^{-7}, 0.638780966e^{-10})$ with indexing starting from 1.

[0064] The RH of each EC model level and at each location within the WRF model domain is then altered by the desired amount that is obtained for example empirically.

Equations 4-7 is then applied in reverse to calculate the new Q_{vap} and this value is then written to a separate intermediate file in the same format as the other intermediate files. The intermediate file with the altered Q_{vap} values is then preferentially utilised by the metgrid program by listing the file prefix last in the namelist. Following these procedures allows a user to alter the Q_{vap} values in the metgrid output. After running metgrid the user is required to run the real program in WRF to provide the input and boundary conditions on the WRF vertical levels. The real program is also responsible for creating the nudging files when grid nudging is turned on. Based on the values pre-set by the user the WRF model can then be run and nudged towards the values from metgrid for the wind components U and V, as well as T and Q_{vap} . The strength of the nudging, the WRF model levels on which to nudge and the duration of the nudging are all controllable by the user through the namelist. Nudging for all of U, V, T and Q_{vap} and for all levels above model level 6 is utilized. The nudging coefficients for U, V and T are set to 0.0002, while that for Q_{vap} it is set at 0.00012.

[0065] To physically address the bias, rather than statistically, the moisture adjustment method described above is used. January 2019, which was one of the worst performing months, is used as the example. The 3D moisture field from the EC model was adjusted by -12% in terms of RH, which in combination with the Q_{vap} nudging coefficient at 0.0012 worked well for January. However, the nudging coefficient can be made stronger (weaker) with a correspondingly smaller (larger) RH % to achieve a similar outcome (not shown). The nudging coefficients are applied through the 54-hour simulation and the Q_{vap} values from EC adjusted in the 15UTC to 15UTC time period corresponding to forecast hours +27 to +51. Adjusting through this time period effectively covers the midnight to midnight local time for the day of concern. It should also be noted that only the values from the EC model grid corresponding to the WRF model grid were altered as it was deemed unnecessary to adjust the whole EC model domain—especially for points in space that would ultimately not be written to the metgrid output files. Figure 33 shows the time series of the optimised cloudy configuration with and without moisture adjustment, as well as the EC model and observed irradiance.

[0066] Figure 34 shows the statistics for January 2019. It can be seen that the combination of grid nudging and the adjustment of the moisture values being nudged towards has a profoundly positive effect on the irradiance profile for the Xinjiang site.

The bias that was observed was almost completely removed (90.5% reduction) and despite the under-estimation by the original WRF model varying from day to day the adjustment of -12% RH appeared suitable for all days in that month. In accordance with the reduction in bias the RMSE also improved by 57.8%.

[0067] Based on the above, the method 2 of forecasting solar irradiance in Figure 1 may alternatively or additionally include adjusting moisture values in the forecast output of the EC model by an adjustment amount corresponding to a window period including the selected day and forecasting with the WRF model is based on the forecast output of the EC model with moisture values adjusted. As described above, the window period may be a calendar month. However, the window period may be shorter or longer than a month.

[0068] Advantageously, improvements in performance are achieved when the WRF model configuration is adjusted based on the type of day—most notably for cloudy days, which had a 15% reduction in RMSE. Adjusting the moisture in the WRF model by means of nudging the model towards adjusted EC model moisture fields is shown to correct any bias from the WRF model. Using January 2019 as an example month it is shown that by drying the EC model by 12% in terms of RH and nudging the WRF model towards these new values produced a 90.5% reduction in bias and a 57.8% reduction in RMSE. A constant 12% drying across the cloudy days of January 2019 is shown to be of a suitable value for most days despite varying amounts of under-prediction across the month.

[0069] Figure 35 shows a block diagram of a computer device according to an embodiment of the present disclosure. The computer device may be configured to implement any of the methods described above. The computer device 600 includes a processing unit (such as a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), and the like) 601, a system memory 604 including a random-access memory (RAM) 602 and a read-only memory (ROM) 603, and a system bus 605 connecting the system memory 604 and the central processing unit 601. The computer device 600 further includes an input/output system (I/O system, basic input/output system) 606 that helps information transmission among various components within a server, and a mass storage device 607 for storing an operating system 613, an application 614 and other program modules 615.

[0070] The I/O system 606 includes a display 608 for displaying information and an input device 609 such as a mouse, a keyboard, and the like, for inputting information by

a user. The display 608 and the input device 609 are both connected to the central processing unit 601 by an input/output controller 610 connected to the system bus 605. The I/O system 606 may further include the input/output controller 610 for receiving and processing input from a plurality of other devices, such as a keyboard, a mouse, or an electronic stylus. Similarly, the input/output controller 610 further provides output to a display screen, a printer or other types of output devices.

[0071] The mass storage device 607 is connected to the central processing unit 601 by a mass storage controller (not shown) connected to the system bus 605. The mass storage device 607 and its associated computer-readable medium provide non-volatile storage for the computer device 600. In other words, the mass storage device 607 may include a computer-readable medium (not shown), such as a hard disk or a drive, or a compact disc read-only memory (CD-ROM).

[0072] The computer-readable medium may include a computer storage medium and a communication medium. The computer storage medium includes volatile and non-volatile medium, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules, or other data. The computer storage medium includes a RAM, a ROM, an erasable programmable read-only memory (EPROM), an electrically erasable programmable read-only memory (EEPROM), a flash memory or other solid-state storage technologies, a CD-ROM, a digital video disc (DVD) or other optical storage, a tape cassette, a magnetic tape, a magnetic disk storage or other magnetic storage devices. Those skilled in the art know that the computer storage medium is not limited to the aforementioned types. The aforementioned system memory 604 and the mass storage device 607 may be collectively referred to as the memory.

[0073] According to the embodiments of the present disclosure, the computer device 600 may also be operated by being connected through a network such as the Internet to a remote computer on the network. That is, the computer device 600 may be connected to the network 612 by a network interface unit 611 connected to the system bus 605, or that is, the computer device 600 may also be connected to other types of networks or remote computer systems (not shown) by using the network interface unit 611.

[0074] The memory further includes a computer program which is stored in the memory and configured to be executed by one or more processors and to cause the one or more processors to implement one or more methods described above.

[0075] In an embodiment of the present disclosure, a non-transitory computer-readable storage medium is further provided, wherein the non-transitory computer-readable storage medium stores thereon a computer program, wherein the computer program, when executed by a processor, causes the processor to implement the methods described above.

[0076] In an exemplary embodiment of the present disclosure, a computer program product is further provided, wherein the computer program product, when executed by a processor, is configured to cause the processor to implement the methods described above.

[0077] Although the present invention is described as implemented in the above described embodiments, it is not to be construed to be limited as such. For example, although it is described that a WRF model is used as a local NWP model, other local NWP models, such as but not limited to, RAMS, MPAS, Unified Model, etc. may also be used.

[0078] As another example, it is described that the EC model is used as a global NWP model, other global NWP models, such as but not limited to, Global Forecast System (GFS), Global Deterministic Prediction System (GDPS), Global Spectral Model (GSM), etc. may also be used.

[0079] As yet another example, it is described that the configurations for the different day-types are optimum. However, it should not be construed to be limited as such. The configuration for each day-type may be any suitable configuration, not necessarily, optimum.

[0080] It should be further appreciated by the person skilled in the art that one or more of the above modifications or improvements, not being mutually exclusive, may be further combined to form yet further embodiments of the present invention.

CLAIMS

1. A method of forecasting solar irradiance for a selected day using a local NWP model based on forecast output of a global NWP model, the method comprising:
 - adjusting moisture values in the forecast output of the global NWP model by an adjustment amount corresponding to a window period including the selected day; and
 - forecasting solar irradiance is based on the forecast output with moisture values adjusted.
2. A method of forecasting solar irradiance according to Claim 1, wherein the window period is a calendar month.
3. A method of forecasting solar irradiance according to Claim 1, wherein adjusting moisture values comprises adjusting water vapour mixing ratios.
4. A method of forecasting solar irradiance according to Claim 3, wherein adjusting water vapour mixing ratios comprises:
 - converting each water vapour mixing ratio into a corresponding relative humidity value;
 - adjusting the relative humidity value by the adjustment amount; and
 - converting the adjusted relative humidity value back to a water vapour mixing ratio.
5. A method of forecasting solar irradiance according to Claim 1, further comprising nudging forecasted values of the local NWP model towards one of adjusted and non-adjusted global NWP model values.
6. A method of forecasting solar irradiance according to Claim 1, further comprising determining an adjustment amount for each window period of a plurality of window periods.

7. A method of forecasting solar irradiance according to Claim 6, wherein each adjustment amount is determined empirically.
8. A method of forecasting solar irradiance according to Claim 6, wherein determining an adjustment amount for each window period only when there is a systematic bias in previously forecasted solar irradiance for the window period.
9. A method of forecasting solar irradiance according to Claim 1, further comprising:
 - writing the adjusted moisture values to a new intermediate file for use by a metgrid program.
10. A computer device for forecasting solar irradiance for a selected day using a local NWP model, the computer device comprising:
 - a processor that is operable to:
 - adjust moisture values in the forecast output of the global NWP model by an adjustment amount corresponding to a window period including the selected day; and
 - forecast solar irradiance is based on the forecast output with moisture values adjusted.
11. A computer device according to Claim 10, wherein the window period is a calendar month.
12. A computer device according to Claim 10, wherein moisture values are adjusted by adjusting water vapour mixing ratios.
13. A computer device according to Claim 12, wherein the processor adjusts water vapour mixing ratios by:
 - converting each water vapour mixing ratio into a corresponding relative humidity value;
 - adjusting the relative humidity value by the adjustment amount; and
 - converting the adjusted relative humidity value back to a water vapour mixing ratio.

14. A computer device according to Claim 10, wherein the processor further nudges forecasted values of the local NWP model towards one of adjusted and non-adjusted global NWP model values.

15. A computer device according to Claim 10, wherein the processor further writes adjusted moisture values to a new intermediate file for use by a metgrid program.