

EARLY ONLINE RELEASE

This is a PDF of a manuscript that has been peer-reviewed and accepted for publication. As the article has not yet been formatted, copy edited or proofread, the final published version may be different from the early online release.

This pre-publication manuscript may be downloaded, distributed and used under the provisions of the Creative Commons Attribution 4.0 International (CC BY 4.0) license. It may be cited using the DOI below.

The DOI for this manuscript is

DOI:10.2151/jmsj.2020-048

J-STAGE Advance published date: June 12th 2020

The final manuscript after publication will replace the preliminary version at the above DOI once it is available.

Intra-day Forecast of Ground Horizontal Irradiance Using Long Short-Term Memory Network (LSTM)

Xiuhong Chen¹, Xianglei Huang^{1*}, Yifan Cai², Haoming Shen³, Jiayue Lu⁴

1. Department of Climate and Space Sciences and Engineering, the University of Michigan, Ann Arbor, Michigan, USA
2. School of Software Engineering, the Shanghai Jiao Tong University, Shanghai, China
3. Department of Industrial and Operations Engineering, the University of Michigan, Ann Arbor, Michigan, USA
4. Department of Computer Science, University of Southern California, Los Angeles, CA, USA

* Corresponding author: Prof. Xianglei Huang, Dept. of Climate and Space Sciences Engineering, the University of Michigan, Ann Arbor, MI 48109-2143, USA.
xianglei@umich.edu, +1-734-936-0491.

Submitted to Journal of the Meteorological Society of Japan on March 8, 2020

The first revision on May 6, 2020

The second revision on May 19, 2020

Abstract

22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40

Accurate forecast of ground horizontal irradiance (GHI) is one of the key issues for power grid managements with large penetration of solar energy. A challenge for solar forecasting is to forecast the solar irradiance with a lead time of 1-8 hours, here termed as intra-day forecast. This study investigated an algorithm using a long short-term memory (LSTM) model to predict the GHI in 1-8 hours. The LSTM model has been applied before for inter-day (>24 hours) solar forecast but never for the intra-day forecast. Four years (2010 - 2013) of observations by the National Renewable Energy Laboratory (NREL) at Golden, Colorado were used to train the model. Observations in 2014 at the same site were used to test the model performance. The results show that, for a 1-4 hour lead time, the LSTM-based model can make predictions of GHIs with root-mean-square-errors (RMSE) ranging from 77 to 143 Wm^{-2} , and normalized RMSEs around 18.4~33.0 %. With 5-minute inputs, the forecast skill of LSTM with respect to smart persistence model is 0.34~0.42, better than random forest forecast (0.27) and the numerical weather forecast (-0.40) made by the Weather Research and Forecasting (WRF) model. The performance levels off beyond 4-hour lead time. The model performs better in fall and winter than in spring and summer, and better under clear-sky conditions than under cloudy conditions. Using adjacent information from the reanalysis as extra inputs can further improve the forecast performance.

Keywords: Atmospheric radiation; Solar energy forecast; Recurrent neural network

41 **1. Introduction**

42 Solar photovoltaic (PV) energy is a major source of renewable energy. Over the last decade
43 it has been increasingly integrated into worldwide electricity grid. Electricity output by solar
44 panels depends on solar irradiance reaching the horizontal surface, a.k.a. global horizontal
45 irradiance (GHI). GHI includes three components: direct solar beam, diffusive solar radiation by
46 atmosphere and clouds, and ground-reflected solar radiation (usually insignificant). The GHI is
47 affected by several environmental factors such as solar irradiance at the top of the atmosphere,
48 cloud fraction, aerosol loading, and relative humidity. As a result, the amount of electricity
49 produced by a solar panel is subjected to fluctuations in such meteorological variables (Bae et al.,
50 2017). This is distinctively different from traditional power generators that are stable over the
51 entire operational period and insensitive to actual weather variations. The effect of such
52 fluctuations becomes significant when PV power generation consists of 5 % of overall power used
53 over a local grid and becomes critical when the PV penetration reaches 20 % (Matsura, 2009). As
54 a result, some traditional generators must be used as a backup (a.k.a. operating reserve) to stabilize
55 the power grid. However, conventional power generation systems (including nuclear) all have their
56 own “warm-up” time, ramp rate, and minimal possible load (IEA, 2014). Therefore, improving the
57 accuracy of solar irradiance forecasting can be critically important for efficient grid management
58 with large penetration of solar energy.

59 Depending on the need for solar energy for different time horizons, different types of solar
60 forecasting techniques can be used. For example, intra-hour forecasting (also known as
61 nowcasting) refers to prediction of GHI from a few minutes up to one hour. It is usually done by
62 simply extrapolating current wind and cloud information with an assumption that cloud property
63 does not change and merely moves with the wind (Bosch and Kleissl, 2013; Chow et al., 2015;

64 Chu et al., 2015; Lorenz et al., 2014; Rana et al., 2016). Thus, only endogenous data are needed
65 for such forecasting. For another example, inter-day forecasting refers to prediction of changes of
66 GHI for one to several days ahead. For such a timescale, the evolution of weather system becomes
67 the leading factor in determining the variations of solar irradiance. As a result, the most useful
68 inter-day forecasts heavily rely on the numerical weather prediction (NWP) (e.g., Alessandrini et
69 al., 2015; Mellit et al., 2014; Almeida et al., 2015). Between the above two ends of time horizon
70 is the intra-day forecast, i.e., solar forecasting with a lead time of 1-8 hours. Many cloud systems
71 can develop, mature, and dissipate within eight hours. The multiscale turbulence nature of cloud
72 systems plays important roles within such timescale. State-of-the-art NWP models still have
73 difficulties simulating such phenomena and, therefore, also have difficulty to well resolve the GHI
74 variations within 8 hours (Cros et al., 2014).

75 The machine-learning approach recently emerges as a promising alternative to do the intra-
76 day forecast since neither simple extrapolation technique nor pure physics-based NWP technique
77 has proved to be fully successful for such type of forecast (Antonanzas et al., 2016; Nonnenmacher
78 and Coimbra, 2014; Miller et al., 2018). A number of algorithms have been used to forecast the
79 GHI over this timescale. For example, David et al. (2018) compared the intraday forecasting of
80 deterministic GHI by three models: coupled autoregressive and dynamical system (CARDS),
81 sequential neural network (NN), and more traditional recursive autoregressive and moving average
82 (ARMA) model. Their results showed that the performance of the ARMA model is comparable to
83 that of NN, but both of them outperform the CARDS.

84 Some studies have employed Long Short-Term Memory networks (LSTMs) to forecast
85 photovoltaic (Gensler et al., 2016), total electron content (Sun et al., 2017), and day-ahead solar
86 irradiance (Qing and Niu, 2018). LSTM is a particular kind of recurrent neural network (RNN;

87 refer to the appendix) introduced by Hochreiter & Schmidhuber (1997). Wang et al. (2018) and
88 Fu et al. (2016) found that a LSTM network can achieve a better performance than the traditional
89 time series prediction methods such as the ARMA and Support Vector Machine Regression. To
90 the best of our knowledge, there has been no use of LSTM for intra-day solar forecasting yet.
91 Therefore, inspired by the success of the LSTM for the inter-day solar forecasting, this study aims
92 at assessing the performance of a LSTM-based model for intra-day solar forecasting. The
93 remainder of this paper is organized as follows. Section 2 describes the structure of the LSTM
94 model used in this study, the observational data used in this study, and the validation metrics used
95 for assessing the performance of the forecast. Section 3 presents the forecasting results and
96 sensitivity tests regarding the choices of input features, temporal resolutions, and some
97 configurations of the LSTM model. Section 4 compares the performance of the LSTM model with
98 that of several other forecast methods. Conclusions and further discussion are given in Section 5.

99 **2. LSTM, Data, and Methodology**

100 **2.1 LSTM-based neural network structure**

101 Figure 1 sketched the structure of LSTM-based time series prediction scheme used in this
102 study. Mathematical description of the LSTM model can be found in the appendix. Unlike normal
103 multi-layer perceptron framework, the LSTM has a special neuron structure called memory cell
104 that can capture short-term memory and keep it for a long time (Graves, 2012). Three gates control
105 the information flow into and out of the memory cell: input gate, output gate, and keep gate (also
106 known as forget gate). If the forecast is made at time t_0 using m sets of observations prior to it, m
107 LSTM models are then needed. The features (i.e. inputs) fed into each LSTM model through the
108 input gate can be GHI only (i.e. one feature only) as well as a set of features including GHI. For
109 our case, the set of features includes GHI, local time, and other observable meteorological

110 variables including relative humidity (RH ; %), total cloud fraction (F_{cld} ; unitless), near-surface
111 temperature (T_s ; K), and solar irradiance at the top of the atmosphere (R_{toa} ; Wm^{-2}), as shown in
112 Figure 1. These meteorological variables all have direct relevance to the solar radiative transfer in
113 the atmosphere. All LSTM models were implemented using Keras neural network library (Chollet,
114 2015), which runs on top of TensorFlow. The output, which has a dimension of `out_d`, from the
115 LSTM model, is then used to as the input to a dense layer to generate forecast of the GHI in the
116 next n hours after time t_0 . The final output for our case is hourly-averaged GHI for the n^{th} hours
117 after time t_0 . Previous studies such as Gensler et al. (2016) and Wang et al. (2018) showed that, in
118 order to do a sequence of forecast, the input for such LSTM framework should be 2-5 times as
119 long as the sequence of the forecast. Our targeted sequence is 8-hour GHI forecast, thus we have
120 tested the length of input from 8 to 56 hours prior to the forecast time. It turns out that using 32-
121 hour observations prior to the forecast time t_0 has the optimal performance. In the following
122 discussions, all results are obtained using 32 hours of features prior to the time of the forecast.

123 The training was deemed finished when the root mean square error decreases to a small value
124 (in our case, $\sim 105 \text{ Wm}^{-2}$ or less) and keeps essentially flat within 50 more training iterations.
125 Coefficients derived from such training were then saved and used to do the 8-hour forecast of GHI.

126 **2.2 Data**

127 Data used in this study was from ground observations at Golden, Colorado (39.74°N ,
128 105.18°W , elevation of 1.829 km) by the National Renewable Energy Laboratory (NREL; Andreas
129 and Stoffel, 1981). Original observations were made every minute. Observed GHI, total cloud
130 fraction, near-surface relative humidity and temperature were used. The GHIs were measured
131 using well-maintained global horizontal pyranometers with a 3 % ~ 5 % measurement uncertainty
132 (Wilcox 2012; Wilcox and Myers, 2008). Four years of observations from 2010 to 2013 were used

133 to train the LSTM network. The observations in the entire year of 2014 were then used for
134 validation. The evaluation of the LSTM forecast performance with respect to different
135 configurations of the network was also carried out using the same sets of training and validation
136 data.

137 **2.3 Metrics for evaluating the model performance: RMSE, nRMSE, and forecast skill**

138 Multiple metrics exist for evaluating the forecasting performance (Orwig et al., 2015). The
139 most commonly used in meteorological forecast is root mean square error (RMSE), which is
140 defined as

$$141 \quad \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\text{GHI}_{\text{predicted},i} - \text{GHI}_{\text{observed},i})^2} \quad (1),$$

142 where N is the number of observations used in the validation. RMSE measures how close the
143 predicted GHIs are to the observed GHIs in the absolute unit of GHI. A closely related metric is
144 normalized RMSE (nRMSE), which is defined as the RMSE divided by the standard deviation of
145 observed GHI, which measures the prediction error with respect to the spread of the observed
146 quantity. nRMSE can be a more useful metric when the value of $\text{GHI}_{\text{observed}}$ varies by several
147 factors or even by orders of magnitude. nRMSE has been used in previous studies about solar
148 forecasting (Deo et al., 2016). Both RMSE and nRMSE will be used here to evaluate the model
149 performance.

150 In addition, we will perform a linear fit of the predicted hourly GHIs with respect to the
151 observed counterparts. The R-square from such linear fit tells the fraction of total variance in
152 observed GHIs that can be explained by a linear function of predicted GHIs. Meanwhile, the slope

153 of the linear fit tells the scaling of such linear function (with 1:1 being the ideal) and the intercept
154 tells the bias of such linear fit.

155 Forecast skill (FS), which is introduced by Coimbra et al. (2013), is widely used as a
156 comparison metric in solar forecast. It is defined as,

$$157 \quad FS = 1 - \frac{RMSE}{RMSE_p} \quad (2),$$

158 Where RMSE is from the model to be assessed. $RMSE_p$ is the RMSE of the persistence model. In
159 this study, the persistence model with clear-sky index is used following Kaur et al. (2016), also
160 known as the smart persistence. The GHI for the n^{th} hour lead time at time t , $GHI(t+n)$, predicted
161 by clear-sky persistence model is defined as,

$$162 \quad GHI(t+n) = k(t+n) \cdot GHI(t+n)_{\text{clr}} \quad (3),$$

$$163 \quad k(t+n) = k(t) = \frac{GHI(t)}{GHI(t)_{\text{clr}}} \quad (4),$$

164 Where $k(t+n)$ and $k(t)$ are the clear-sky index, they are the same for the current time and for the
165 n^{th} hour lead time. $GHI(t)_{\text{clr}}$ and $GHI(t+n)_{\text{clr}}$ are the clear-sky solar radiations at surface for the
166 time t and $t+n$. They are from solar radiation data (SoDa; Lefèvre et al., 2013). A forecast skill
167 of 1.0 indicates a perfect forecasting. A forecast skill of 0 indicates an exact same performance as
168 the persistence model. And a forecast skill of <0 indicates even worse performance than the
169 persistence model.

170 **3. Results**

171 **3.1 LSTM prediction Performance**

172 First, 32 hours of high temporal resolution data (5 minutes per data point) was used as input
173 to the LSTM model with `out_d=100` to forecast GHIs in the next eight hours with also a temporal

174 resolution of five minutes. In other words, 384 sequential inputs (i.e. 32-hour observations) were
175 used to predict the next 96 sequential outputs (i.e. GHI forecast for the next eight hours). The input
176 data vectors have six features including GHIs, time, total cloud fraction, relative humidity, near
177 surface temperature, and solar irradiance at the top of atmosphere (TOA). Only daytime forecast
178 is examined further here, as the nighttime GHI is zero and has no relevance to solar power
179 generation.

180 Figure 2 shows the scatter plots between the predicated hourly-averaged GHIs and observed
181 counterparts for a 1-hour lead time (also referred to as 1-hour forecast horizon). Results from four
182 seasons are plotted separately. The largest deviation from 1:1 line is seen in spring, which has a
183 slope of 1.09. All slopes are larger than 1 and, correspondingly, the intercepts are all negative,
184 ranging from -2 Wm^{-2} in winter to -20.5 Wm^{-2} in spring. The R-squares range from 0.88 in the
185 summer to 0.93 and 0.92 in the fall and winter seasons, respectively. In general, the prediction
186 performance is the best in winter and the poorest in summer. This is consistent with our
187 understanding of the weather system. In winter large-scale, long-lived weather systems affect the
188 cloud evolutions the most. For example, when a synoptic frontal system passes the site, it usually
189 is featured with overcast sky over a few days. Then the high-pressure system comes with entire
190 clear sky or little cloud coverage over a few days. In summer, however, large-scale weather system
191 is less organized; diurnal cycle featured with local, short-lived convective storms dominates the
192 cloud variations and thus the GHI is less predictable than in the winter. This can also be seen from
193 the scatter plots in Figure 2: the linear model fit the wintertime data much better than the
194 summertime data.

195 The RMSE, nRMSE, and linear regression outcomes (slope and intercept) for a lead time
196 from one hour to eight hours are summarized in Figure 3. For annual average, the RMSE (nRMSE)

197 increases from $\sim 77 \text{ Wm}^{-2}$ ($\sim 18\%$) at one-hour lead time to $\sim 143 \text{ Wm}^{-2}$ ($\sim 33\%$) at four-hour lead
198 time. Both RMSE and nRMSE suggest that the forecast performance levels off beyond four hours.
199 As expected, the wintertime RMSE is always the smallest among all seasons and for all lead times.
200 Except for winter, the regression slopes (intercepts) for other seasons decrease (increase) with the
201 increase of lead time, which can be explained by the fact that randomness increases with respect
202 to the lead time (e.g. a convective storm or shower happens after forecast time t_0 and lasts for only
203 a hour or two), and such randomness has little dependence on what happens before t_0 . As explained
204 above, winter weather is mostly governed by large-scale synoptic system that could last several
205 days, which can account for the seasonal contrast here.

206 The dependence of forecast performance with respect to cloud fraction is also examined.
207 Figure 4 shows the performance metrics with respect to four groups of observed cloud fraction.
208 The RMSE for the group with cloud fraction $< 1\%$ is significantly less than those for the rest three
209 groups. The RMSE for such clear or nearly clear sky situation mainly rises from the aerosol
210 absorption and scattering of sunlight and the water vapor absorption of solar radiation. While the
211 input features such as surface temperature and relative humidity are related to water vapor, due to
212 measurement limitation our input features do not include information about aerosols. When cloud
213 fraction is less than 95 %, the RMSE and nRMSE both level off after the 4th hour. However, this
214 is not the case for overcast or nearly overcast situation, where both RMSE and nRMSE keep
215 increasing with forecast horizon. The slope of regression also keeps dropping with increase of
216 forecast horizon, dropping to as low as 0.6 for the 8th hour, but the intercept of regression is always
217 less than 10 Wm^{-2} and is nearly zero for the 8th hour. These highlight the challenge of solar
218 forecasting for overcast skies. From the meteorological perspective, overcast sky implies small
219 GHI and even a small amount of reduction in cloud fraction can quickly increase the amount of

220 solar irradiance reaching the surface. Within eight hours, an overcast sky at t_0 can evolve to any
221 possibilities (e.g., still overcast, completely clear sky, and partly cloudy), which explains at least
222 partly why the forecast performance for this group is not as good as for the rest three groups and
223 why it has no apparent level-off forecast horizon.

224 **3.2 Sensitivity of the LSTM performance to different configurations**

225 We further examine forecast performance with respect to the inputs at different temporal
226 resolutions, to the different choices of input features, and to a key hyperparameter, the number of
227 the last LSTM outputs (out_d). The goal is to estimate an optimal configuration for the LSTM
228 forecast algorithm in terms of computational speed and accuracy. RMSE and nRMSE will be used
229 in the assessment of different configurations, as they are more useful than the linear regression
230 statistics for such purposes (Orwig et al., 2015).

231 **3.2.1. Temporal resolutions**

232 Five temporal resolutions for the input features are tested, i.e., 5-minute, 15-minute, 20-
233 minute, 30-minute, and 1-hour inputs. Other configurations are identical to what is used in Section
234 3.1. For each temporal resolution, the LSTM network is trained separately. The performances are
235 shown in Figure 5. After the 4th hour, the RMSE and nRMSE both show little dependence on the
236 temporal resolutions of the input features. Within four hours, a finer temporal resolution can
237 improve the performance but the dependence is stepwise instead of linear. For example, the 5-
238 minute and 15-minute inputs have nearly identical performance, so do the 20-minute and half-hour
239 inputs. The improvements from one-hour to half-hour resolution, in terms of RMSE and nRMSE,
240 are $\sim 20\%$ while from half-hour to five-minute resolution the improvement is only $\sim 10\%$. On the
241 same computer platform, it takes 5, 22, 50, 88, and 790 seconds to finish one iteration of training
242 for the inputs with 1-hour, 30-minute, 20-minute, 15-minute and five-minute resolutions,

243 respectively. The prediction takes 37, 134, 306, 480 and 4848 seconds for the inputs with 1-hour,
244 30-minute, 20-minute, 15-minute and five-minute resolution, respectively. The computational cost
245 from 15-minute to five-minute resolution increases by a factor of ~ 10 while the forecast
246 performance is comparable. Given the RMSE dependence and the computing cost, inputs with 15-
247 minute resolution are recommended among all temporal resolutions examined here.

248 **3.2.2 Choices of input features**

249 Among six input features, the most important one is observed GHI. It is meaningful to look
250 a LSTM network using the GHI as the only input feature and to compare its performance with the
251 counterpart that uses all six features as input to train the LSTM network. Figure 6 summarizes the
252 comparisons of performance between six features and only the GHI feature. For the 1-hour inputs,
253 though the 6-feature prediction always performs better than the one-feature prediction, the
254 difference is small for the first two hours ($\sim 2\%$) and then increases with the forecast horizon to
255 $\sim 5\%$. For the 15-minute inputs, the difference is only 1% or less. On the same platform, one
256 iteration in training of 6-feature and one-feature networks takes nearly the same amount of time.
257 Given no difference in training, six features should be advocated instead of one feature. But if the
258 data availability is an issue, GHI observation with high temporal resolution (15-minute or even
259 higher) alone can be used for such forecast.

260 **3.2.3 Choice of out_d**

261 Five configurations of out_d are tested. The performances are shown in Figure 7. For 1-hour
262 inputs (Figure 7a and 7c), when out_d varies from 50 to 100, the performance is essentially the
263 same. When out_d increases from 10 to 50, the performance is significantly improved for the first-
264 hour prediction. For the case of 15-minute (Figure 7b and 7d), however, changing out_d from 10

265 to 100 nearly leads to no improvement in the performance for all forecast horizons. The trade-off
266 between out_d and the temporal resolution of the inputs can be clearly seen by contrasting left
267 panels and right panels in Figure 7. For a 1-5 hour lead time, the performances with any out_d
268 using 15-minute inputs are better than the performances with out_d =100 and 1-hour inputs.
269 Beyond the 5th hour, however, the performance using 15-minute inputs becomes comparable to
270 the performance using 1-hour inputs with out_d \geq 50. The time to train a LSTM network with
271 out_d = 100 is 2.5 times as much as the time needed to train a LSTM network with out_d = 10.
272 Based on these facts, 15-minute inputs with out_d =10 is an affordable option with performance
273 no worse than any other configurations examined here.

274 **3.2.4 With additional neighboring features from reanalysis**

275 This section is to test whether the features from adjacent locations can help improve the
276 forecast. Given the actual GHI observations at neighboring locations to this NREL site are not
277 available, we use hourly GHIs from the latest European Centre for Medium-Range Weather
278 Forecasts (ECMWF) ERA5 reanalysis (C3S, 2017) over a 0.75-degree by 0.75-degree domain
279 centered at the grid box encompassing NREL location. The hourly GHIs from ERA5 is at a
280 resolution of 0.25-degree by 0.25-degree. The ERA5 GHIs on 3-by-3 grids are thus used as input
281 features to the LSTM to predict the GHI at NREL, together with other features used in the previous
282 subsections. Figure 8 shows the differences between with and without GHIs from the ERA5
283 reanalysis, for RMSE and nRMSE, respectively. The differences are negative up to a 7-hour lead
284 time, with maximum difference being 5 Wm⁻² for the RMSE and 1% for the nRMSE. Therefore,
285 using neighboring GHI features near the observed site can help the prediction performance.

286 **4. The forecast skill and comparisons with other forecast methods**

287 Using the forecast skill (FS) defined in subsection 2.3 as a metric, we compare the
288 performance of LSTM with that of the random forest method (Breiman, 2001) and of numerical
289 weather prediction made by the NCAR's Weather Research and Forecasting version 3.9 (WRF;
290 Skamarock et al., 2008), as well as the FS from a physical method published by Miller et al. (2018).
291 The same 5-minute data used for the LSTM model training is used for the random-forest training.
292 The WRF model forecast domain is centered at NREL Golden site with an outside domain of 192
293 km by 192 km (a 3-km grid on each direction) nested with an inner domain of 64 km by 64 km (1-
294 km grid on each direction). The model used NCEP operational Global Forecast System (GFS)
295 analysis as initializations and lateral boundary conditions for a 8-hourly daytime forecast every
296 day in 2015, and then the GHI forecast were evaluated against observation.

297 The overall comparisons are summarized in Table 1. The LSTM algorithm outperforms the
298 random forest, with $\sim 2\%$ reduction for nRMSE and 8-10 Wm^{-2} reduction in RMSE. Compared to
299 numerical weather forecast by the WRF model, the LSTM algorithm can reduce RMSE by 54~68
300 Wm^{-2} and nRMSE by $\sim 12\%$. Miller et al. (2018) used a NWP model in combination with satellite-
301 based cloud advection techniques to predict GHIs with a lead time of 1-3 hours. They evaluated
302 their method at four different locations including the Table Mountain in Boulder, Colorado
303 (40.125° N, 105.237° W), which is close to the NREL site. They found that, for this site over a
304 three-year period, the 1-hour and 3-hour forecast RMSEs are 194.7 and 237.9 Wm^{-2} , respectively.
305 The forecast skill is negative. These findings are comparable to our WRF model results (Table 1).
306 In contrast, the 3-hour and 1-hour forecast skills of LSTM are 0.22~0.33 for hourly inputs and
307 0.37~0.42 for 5-minute inputs, better than those of NWP forecast. Forecast skill for all 1-8 hour
308 lead time is shown in Figure 9. Overall, the forecast skill in 2014 ranges from 0.22 to 0.39 for
309 hourly inputs and 0.34 to 0.42 for 5-minute inputs.

310 5. Conclusion and Discussion

311 This study designed a LSTM-based model to perform intra-day GHI forecast and evaluated
312 its performance using observations from a single station at Golden, Colorado by NREL. Four years
313 of observations from 2010 to 2013 were used to train the model, and the evaluation is based on the
314 prediction for the entire year of 2014. The normalized RMSE is ~20 % for the 1st-hour prediction.
315 The performance is better in the winter (RMSE = 58.0 Wm⁻²) than in the summer (RMSE = 109.1
316 Wm⁻²), which is consistent with our physical understanding of the traits of different weather
317 systems governing the GHI variations in different seasons. The forecast performance levels off
318 after the fourth hour with a nRMSE around 30%. The forecast skill to smart persistence model is
319 0.22 to 0.39 for hourly inputs and 0.34 to 0.42 for 5-minute inputs. Such performance is better than
320 that of forecast based on random forest method (0.27), and is much better than that of numerical
321 weather forecast using the WRF model (-0.40). The performance within first four hours show
322 dependence on the temporal resolution of the input features: inputs with high temporal resolution,
323 in general, perform better than those with low temporal resolution but the performance shows step-
324 wise dependence with temporal resolution. Using multiple features with physical relevance to GHI
325 in the context of atmospheric radiation outperforms a single GHI feature, especially for forecast
326 beyond two hours. But there is a trade-off between the temporal resolution of the inputs and the
327 number of input features, as well as a trade-off between the temporal resolution of the inputs and
328 the choice of a key hyperparameter, i.e., the number of outputs from the last cell of the LSTM
329 (out_d). With computational cost and accuracy both considered, an optimal LSTM configuration
330 for the intra-day GHI forecasting can be 15-minute GHI as the only input to LSTMs with out_d
331 =10. We also found that using the GHI features from neighbor locations (such as GHIs derived
332 from the ERA5 reanalysis) can improve the LSTM model forecast performance.

333 This study explores the potential of LSTM in solar forecasting with the use of only
334 indigenous observations. Limited by observational data, some key physical factors such as aerosol
335 loadings were not available for this study. In principle introducing such features that are directly
336 related to GHI can have promise to further improve the performance of machine-learning based
337 algorithm, which is one pathway that we are interested in pursuing further. Meanwhile, introducing
338 features from other data sources, such as simultaneous satellite observations and information about
339 large-scale weather patterns from NOAA National Weather Service, also have a chance to improve
340 the performance of such algorithms further. If the LSTM, as one type of recurrent neural network
341 (RNN), can be successfully employed for solar forecasting, it is encouraging to further explore
342 other RNNs in such forecast applications.

343 **Acknowledgments**

344 We want to thank two anonymous reviewers for their comments and suggestions. We also
345 thank the Editorial Committee members for their comments and careful proofreads. The National
346 Renewable Energy Laboratory (NREL) ground measurements are downloadable at
347 <https://midcdmz.nrel.gov/apps/go2url.pl?site=BMS>. SoDa data is downloaded from
348 <http://www.soda-pro.com/web-services/radiation/cams-mcclear>. GFS forecast and analysis data is
349 downloaded from <https://rda.ucar.edu/datasets/ds084.1/>. Upon acceptance, the code used in this
350 study will be available at https://github.com/xiuhongchen/Solar_forecast_LSTM. This work was
351 supported by Office of Research at the University of Michigan.

352

353

Appendix

A description of LSTM neural network cell

Details of LSTM neural network can be found in Olah (2015) and Sun et al. (2017). Below is a brief summary of the essence of LSTM. Figure A1 shows the typical LSTM neural network cell. The data flow in and out of three type of gates, namely $\mathbf{h}(t-1)$ and $\mathbf{x}(t)$ as two input gates, $\mathbf{C}(t)$ as a keep gate, and $\mathbf{h}(t)$ as an output gate. $\mathbf{h}(t-1)$ contains outputs from the previous LSTM cell state and $\mathbf{x}(t)$ are the current inputs directly from the input layer. $\mathbf{C}(t)$ and $\mathbf{h}(t)$ are expressed as :

$$\mathbf{C}(t) = \mathbf{f}_t * \mathbf{C}(t - 1) + \mathbf{i}_t * \hat{\mathbf{C}}_t \quad (\text{A.1}),$$

$$\mathbf{h}(t) = \mathbf{o}_t * \tanh[\mathbf{C}(t)] \quad (\text{A.2}),$$

where $*$ stands for the scalar product of two vectors or matrices, and

$$\mathbf{f}_t = \sigma(\mathbf{W}_{f,x} \cdot \mathbf{x}(t) + \mathbf{W}_{f,h} \cdot \mathbf{h}(t - 1) + b_f) \quad (\text{A.3}),$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_{i,x} \cdot \mathbf{x}(t) + \mathbf{W}_{i,h} \cdot \mathbf{h}(t - 1) + b_i) \quad (\text{A.4}),$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_{o,x} \cdot \mathbf{x}(t) + \mathbf{W}_{o,h} \cdot \mathbf{h}(t - 1) + b_o) \quad (\text{A.5}),$$

$$\hat{\mathbf{C}}_t = \tanh(\mathbf{W}_{\hat{c},x} \cdot \mathbf{x}(t) + \mathbf{W}_{\hat{c},h} \cdot \mathbf{h}(t - 1) + b_{\hat{c}}) \quad (\text{A.6}),$$

where σ denotes the sigmoid function and \tanh denotes hyperbolic tangent function. The parameters to be determined by training are weight matrices $\mathbf{W}_{f,x}$, $\mathbf{W}_{f,h}$, $\mathbf{W}_{i,x}$, $\mathbf{W}_{i,h}$, $\mathbf{W}_{o,x}$, $\mathbf{W}_{o,h}$, $\mathbf{W}_{\hat{c},x}$, $\mathbf{W}_{\hat{c},h}$ and bias vectors b_f , b_i , b_o and $b_{\hat{c}}$.

Based on above formulas, the interpretation of each LSTM cell can be given as follow. \mathbf{f}_t is between 0 and 1, with one being completely keeping the previous cell state, $\mathbf{C}(t-1)$, and zero being

373 completely discarding it. $\hat{\mathbf{C}}_t$ are new candidate values derived from tanh function with input of
374 both outputs from the previous LSTM cell, $\mathbf{h}(t-1)$, and the current input, $\mathbf{x}(t)$. \mathbf{i}_t determines to what
375 extent $\hat{\mathbf{C}}_t$ is used for updating the current state $\mathbf{C}(t)$. \mathbf{o}_t decides the fraction of $\tanh(\mathbf{C}(t))$ to be used
376 for computing the current output $\mathbf{h}(t)$.

377 **References:**

378 Alessandrini, S., L. Delle Monache, S. Sperati, G. Cervone, 2015: An analog ensemble for short-
379 term probabilistic solar power forecast. *Appl. Energy*, **157**, 95–110.

380 Almeida, M. P., O. Perpignan, L. Narvarde, 2015: PV power forecasting using a nonparametric PV
381 model. *Sol. Energy*, **115**, 354–368.

382 Andreas, A. and T. Stoffel, 1981: NREL Solar Radiation Research Laboratory (SRRL): Baseline
383 Measurement System (BMS); Golden, Colorado (Data); NREL Report No. DA-5500-56488.
384 <http://dx.doi.org/10.5439/1052221>.

385 Antonanzas, J., N. Osorio, R. Escobar, R. Urraca, F.J. Martinez-de-Pison, F. Antonanzas-Torres,
386 2016: Review of photovoltaic power forecasting. *Sol. Energy*, **136**, 78-111.

387 Bae, K.Y. and Co-authors, 2017: Hourly Solar Irradiance Prediction Based on Support Vector
388 Machine and Its Error Analysis. *IEEE Trans. Power Systems*, **32**, 935-945.

389 Bosch, J. and J. Kleissl, 2013: Cloud motion vectors from a network of ground sensors in a solar
390 power plant. *Sol. Energy*, **95**, 13–20.

391 Breiman, L., 2001: Random forests, *Mach. Learn.* **45**, 5-32.

392 Chollet, F., 2015: Keras. Accessed 6 November 2018, <https://github.com/fchollet/keras>.

393 Chow, C., S. Belongie, J. Kleissl, 2015: Cloud motion and stability estimation for intra-hour solar

394 forecasting. *Sol. Energy*, **115**, 645–655.

395 Chu, Y., B. Urguhart, S. Gohari, H. Pedro, J. Kleissl, C. Coimbra, 2015: Short-term reforecasting
396 of power output from a 48MWe solar PV plant. *Sol. Energy*, **112**, 68–77.

397 Coimbra, C.F.M., J. Kleissl, R. Marquez, 2013: Overview of solar forecasting methods and a
398 metric for accuracy evaluation. In *Solar Energy Forecasting and Resource Assessment*, 1st ed.;
399 Kleissl, J., Ed.; Elsevier Academic Press: Boston, MA, USA, 2013; Chapter 8, pp. 171–194,
400 ISBN 9780123971777.

401 Copernicus Climate Change Service (C3S), 2017: ERA5: Fifth generation of ECMWF
402 atmospheric reanalyses of the global climate . Copernicus Climate Change Service Climate
403 Data Store (CDS). Accessed 20 November 2019,
404 <https://cds.climate.copernicus.eu/cdsapp#!/home>.

405 Cros, S., S. Jolivet, N. Sébastien, N. Schmutz, 2014: Using Satellite Data for Intraday Solar Energy
406 Forecasts: Benchmarks of Two Cloud Trajectory Retrieval Techniques. 29th European
407 Photovoltaic Solar Energy Conference and Exhibition, 2220 - 2224,
408 <https://doi.org/10.4229/EUPVSEC20142014-5AO.8.4>.

409 David, M., L. Mazorra, P. Lauret, 2018: Comparison of intraday probabilistic forecasting of solar
410 irradiance using only endogenous data. *Int. J. Forecast*, **34**, 529–547.

411 Deo, R.C., X.H. Wen, F. Qi, 2016: A wavelet-coupled support vector machine model for
412 forecasting global incident solar radiation using limited meteorological dataset. *Appl. Energy*,
413 **168**, 568–593.

414 Fu, R., Z. Zhang, L. Li, 2016: Using LSTM and GRU neural network methods for traffic flow
415 prediction. In *Proceedings of the Youth Academic Annual Conference of Chinese Association*

416 *of Automation (YAC)*, Wuhan, China, 11–13 November, pp. 324–328.

417 Gensler, A., J. Henze, B. Sick, and N. Raabe, 2016: Deep learning for solar power forecasting-An
418 approach using AutoEncoder and LSTM Neural Networks. In *Systems, Man, and Cybernetics*
419 *(SMC)*, 2016 IEEE International Conference on (pp. 002858-002865).

420 Graves, A., 2012: Supervised sequence labelling with recurrent neural network. *Studies in*
421 *Computational Intelligence*, Chapter 4: Long short-term memory, P31-38. Springer.

422 Hochreiter, S. and J. Schmidhuber, 1997: Long short-term memory. *Neural Computation*, **9(8)**,
423 1735-1780.

424 IEA (International Energy Agency), 2014: The power of transformation - Wind, Sun and the
425 Economics of Flexible Power Systems. Edition, 238pp, Paris, OECD/IEA.

426 Kaur, A., L. Nonnenmacher, H.T. Pedro, and C.F. Coimbra, 2016: Benefits of Solar Forecasting
427 for Energy Imbalance Markets. *Renew. Energy*, **86**, 819–830.

428 Lefèvre, M., A. Oumbe, P. Blanc, B. Espinar, B. Gschwind, Z. Qu, L.Wald, M. Schroedter-
429 Homscheidt, C. Hoyer-Klick, A. Arola, A. Benedetti, J. W. Kaiser, and J.-J. Morcrette, 2013:
430 McClear: a new model estimating downwelling solar radiation at ground level in clear-sky
431 conditions. *Atmos. Meas. Tech.*, **6**, 2403-2418.

432 Lorenz, E. and Co-authors, 2014: PV power predictions on different spatial and temporal scales
433 integrating PV measurements, satellite data and numerical weather predictions. In: *29th*
434 *EUPVSEC*, Amsterdam, Netherland.

435 Matura, M., 2009: Island Breezes: Renewable Energy Integration from a Hawaiian Perspective.
436 *IEEE Power Energy Mag.*, **7(6)**, 59-64.

437 Mellit, A., A. Massi Pavan, V. Lughi, 2014: Short-term forecasting of power production in a large-

438 scale photovoltaic plant. *Sol. Energy*, **105**, 401–413.

439 Miller, S.D., A.R. Matthew, J.M. Haynes, M. Sengupta, 2018: Short-term solar irradiance
440 forecasting via satellite/model coupling. *Sol. Energy*, **168**, 102-117.

441 Nonnenmacher, L. and C. F. M. Coimbra, 2014: Streamline-based method for intra-day solar
442 forecasting through remote sensing. *Sol. Energy*, **108**, 447–459.

443 Olah, C., 2015. Accessed 8 February 2019, [http://colah.github.io/posts/2015-08-Understanding-](http://colah.github.io/posts/2015-08-Understanding-LSTMs/)
444 [LSTMs/](http://colah.github.io/posts/2015-08-Understanding-LSTMs/).

445 Orwig, K.D., M.L. Ahlstrom, V. Banunarayanan, J. Sharp, J.M. Wiczak, J. Freedman, S.E. Haupt,
446 J. Cline, O. Bartholomy, H.F. Hamann, B-M Hodge, C. Finley, D. Nakafuji, J.L. Peterson, D.
447 Maggio, M. Marquis, 2015: Recent Trends in Variable Generation Forecasting and Its Value to
448 the Power System. *IEEE Trans. Sustain. Energy*, **6 (3)**, 924–933.

449 Qing, X., Y. Niu, 2018: Hourly day-ahead solar irradiance prediction using weather forecasts by
450 lstm. *Energy*, **148**, 461–468.

451 Rana, M., I. Koprinska, V. Agelidis, 2016: Univariate and multivariate methods for very short-
452 term solar photovoltaic power forecasting. *Energy Convers. Manage*, **121**, 380–390.

453 Skamarock, W., J. Klemp, J. Dudhia, et al., 2008: A description of the advanced research wrf
454 version 3. Technical Note NCAR/TN-475+STR. 125 pp. University Corporation for
455 Atmospheric Research. <http://dx.doi.org/10.5065/D68S4MVH>.

456 Sun, W.Q., L. Xu, X. Huang, W.Q. Zhang, T.J. Yuan, Z. Chen, Y.H. Yan, 2017: Forecasting of
457 ionospheric vertical total electron content (TEC) using LSTM networks. In: *Proc. Int. Conf.*
458 *Mach. Learn. Cybern. (ICMLC)*, pp. 340–344.

459 Wang, Y., W. Liao, Y. Chang, 2018: Gated Recurrent Unit Network-Based Short-Term
460 Photovoltaic Forecasting. *Energies*, **11**, 2163.

461 Wilcox, S.M., 2012: National Solar Radiation Database 1991–2010 Update: User’s Manual
462 (Technical Report NREL/TP-5500-54824). Golden, CO: National Renewable Energy
463 Laboratory. <http://www.nrel.gov/docs/fy12osti/54824.pdf>.

464 Wilcox, S.M., and D. Myers, 2008: Evaluation of Radiometers in Full-Time Use at the National
465 Renewable Energy Laboratory Solar Radiation Research Laboratory (Technical Report
466 NREL/TP-550-44627). Golden, CO: National Renewable Energy Laboratory.
467 <http://www.nrel.gov/docs/fy09osti/44627.pdf>.

468

469 Table 1 Comparison of the performance of the LSTM with the performance by other machine
 470 learning model, Random Forest algorithms and by other numerical models. RMSE refers to root-
 471 mean-square error and nRMSE refers to normalized RMSE as defined in the text. Forecast skill
 472 for different years is with respect to the results of the persistence model in the corresponding
 473 year(s). Note these results are the statistics of daytime GHIs. Random Forest is done only for 1
 474 hour lead time.

	Time period	RMSE (Wm^{-2})		nRMSE(%)		Forecast skill	
		1 hour	3 hour	1 hour	3 hour	1 hour	3 hour
LSTM(5min-input)	2014	78.2	134.2	17.4	29.9	0.42	0.37
LSTM(1hour-input) + ERA5	2014	106.0	140.8	23.2	30.8	0.22	0.33
	2015	100.6	129.8	22.0	28.5	0.20	0.33
Random Forest (5min-input)	2014	98.7	NA	21.0	NA	0.27	NA
WRF	2015	176.3	197.5	35.9	43.0	-0.40	-0.01
Satellite+model (Miller et al., 2018)	2014-2016	194.7	237.9	24.9	28.8	-0.57	-0.21
Smart persistence	2014	135.6	211.4	28.7	41.7		
	2015	126.1	194.9	26.8	38.6		
	2014-2016*	123.9	196.0	26.0	38.3		

475 *The persistence results for these 3 years are calculated using SURFRAD GHIs and SoDa clear-
 476 sky irradiances at Table mountain, CO (40.125° N, 105.237° W). SURFRAD GHIs are
 477 downloaded at <ftp://aftp.cmdl.noaa.gov/data/radiation/surfrad/>.

478

Figure Captions

479 Figure 1. LSTM-based time series prediction framework. The input feature(s) is denoted as x and
480 $(m+1)$ time steps in the past are used for input to predict m time steps in the future. Three gates
481 control the information flow in and out of each LSTM cell, namely input gate, output gate and
482 keep gate.

483 Figure 2. Scatter plot of the predicted versus observed hourly-averaged GHI. The prediction is for
484 one-hour lead time. The entire year of 2014 data were used. Solid black lines are the 1:1 reference
485 lines. Linear regression results are shown as red line. Each panel is for one season.

486 Figure 3. The LSTM performance for hourly averaged GHI prediction with different lead time
487 from one to eight hours. (a) regression slope, (b) intercept, (c) RMSE, and (d) nRMSE. Results
488 from different seasons are shown in different colors. The performances for the entire year are
489 shown in black curves.

490 Figure 4. Performance of the LSTM forecast composite with respect to four groups of cloud
491 fraction at forecast time zero: cloud fraction $<1\%$ in red, $1-50\%$ in red, $50-95\%$ in blue and $>95\%$
492 in green. Four panels are the regression slope (a), regression intercept (b), RMSE (c) and
493 nRMSE (d).

494 Figure 5. Sensitivity of RMSE (a) and nRMSE (b) to the temporal resolution of the inputs. The
495 same LSTM configuration used for Figures 2-4 is used here.

496 Figure 6. Sensitivity of RMSE (a) and nRMSE (b) to number of features used as the inputs. Results
497 predicted using 1-hour inputs and 15-minute inputs are in solid and dashed lines, respectively. One
498 feature is to use GHI only. Six features include time, GHI, total cloud fraction, solar irradiance at

499 TOA, relative humidity, and air temperature. The same LSTM configuration used for Figures 2-4
500 is used here.

501 Figure 7. Sensitivity of RMSE and nRMSE to the number of LSTM output (out_d). Left column:
502 1-hour inputs are used. Right column: 15-minute inputs are used. Upper row is for RMSE and
503 lower row is for nRMSE. Results from different out_ds are shown in different colors, as labeled
504 in (a) for all the panels.

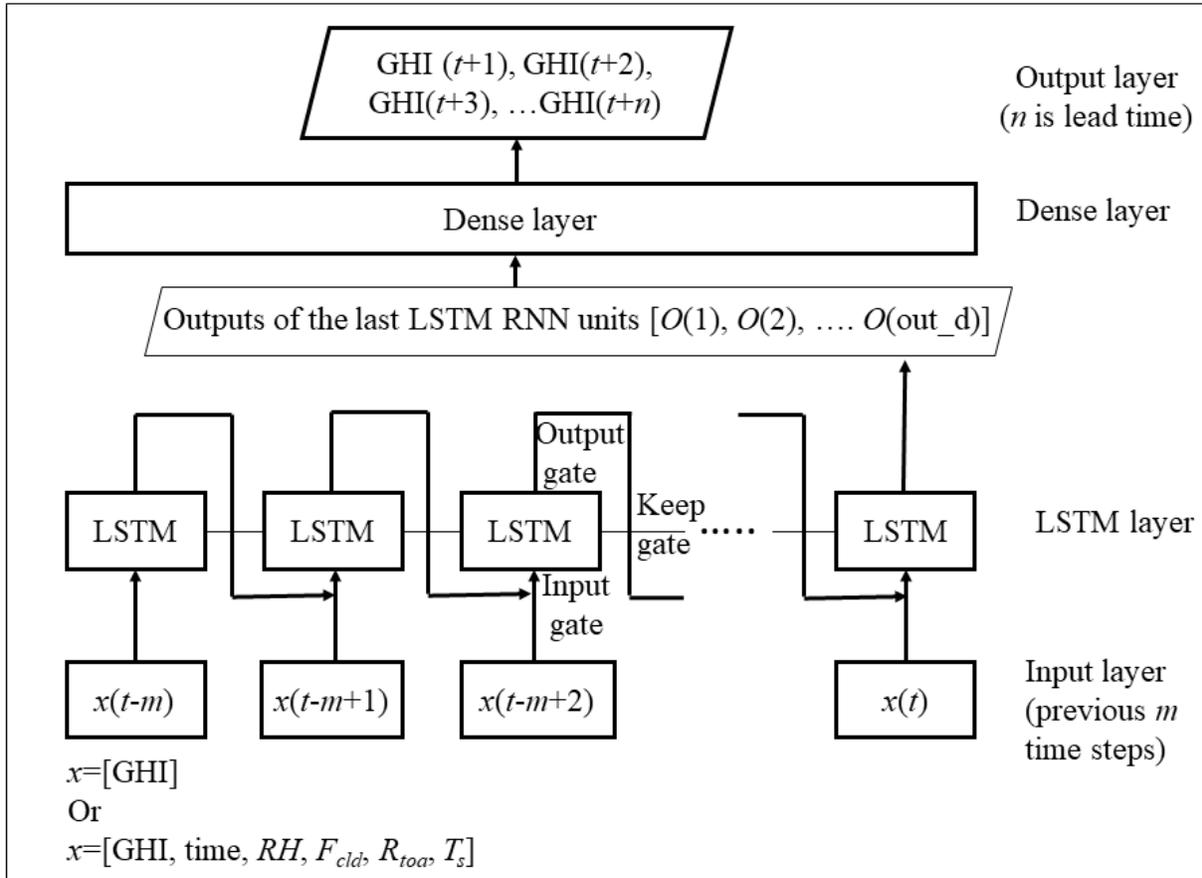
505 Figure 8. Difference in RMSE and nRMSE between with and without hourly ERA5 reanalysis
506 used. Hourly inputs are used, other configurations are the same as in Figures 2-4. For black line,
507 data in 2010-2013 are used to predict GHIs in 2014. For red line, data in 2011-2014 are used to
508 predict GHIs in 2015.

509 Figure 9. Forecast skill in 2014 compared to smart persistence model.

510 Figure A.1 Sketch of the structure of LSTM neural network cell based on illustration from Olah
511 (2015) and Sun et al. (2017).

512

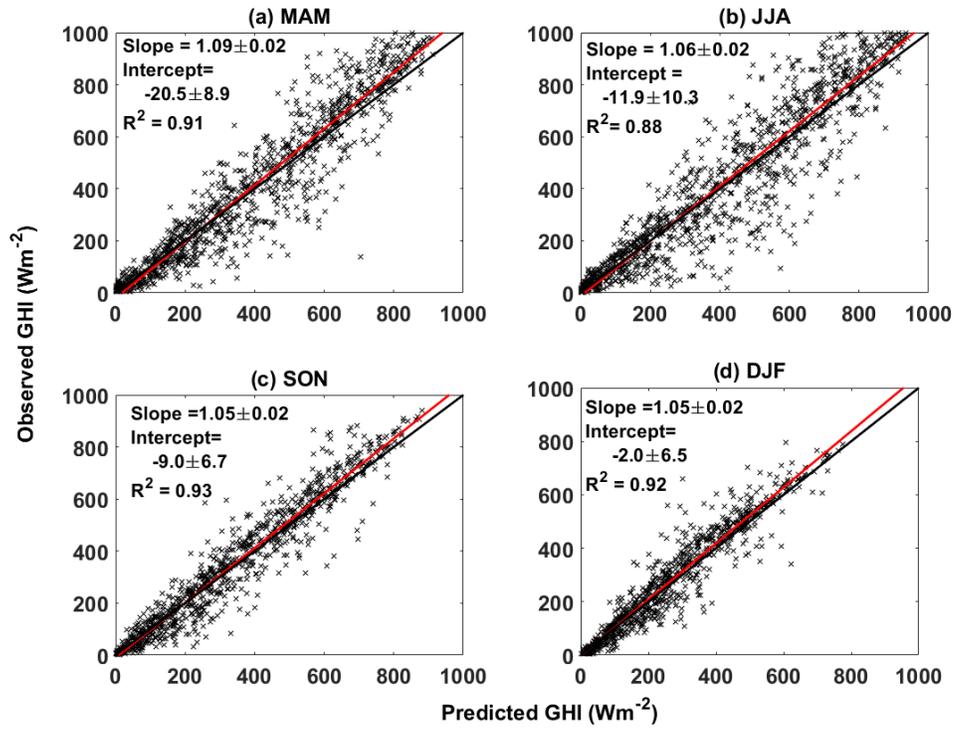
513



514

515 Figure 1. LSTM-based time series prediction framework. The input feature(s) is denoted as x and
 516 $(m+1)$ time steps in the past are used for input to predict m time steps in the future. Three gates
 517 control the information flow in and out of each LSTM cell, namely input gate, output gate and
 518 keep gate.

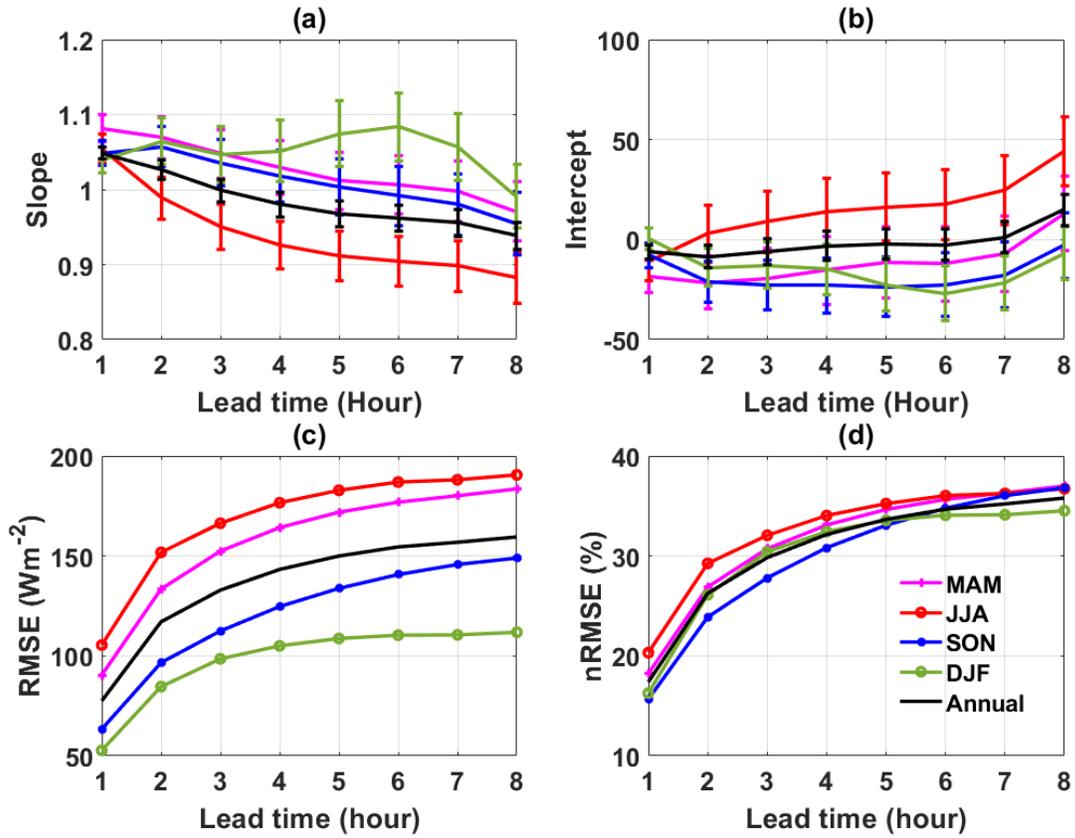
519



520

521 Figure 2. Scatter plot of the predicted versus observed hourly-averaged GHI. The prediction is for
 522 one-hour lead time. The entire year of 2014 data were used. Solid black lines are the 1:1 reference
 523 lines. Linear regression results are shown as red line. Each panel is for one season.

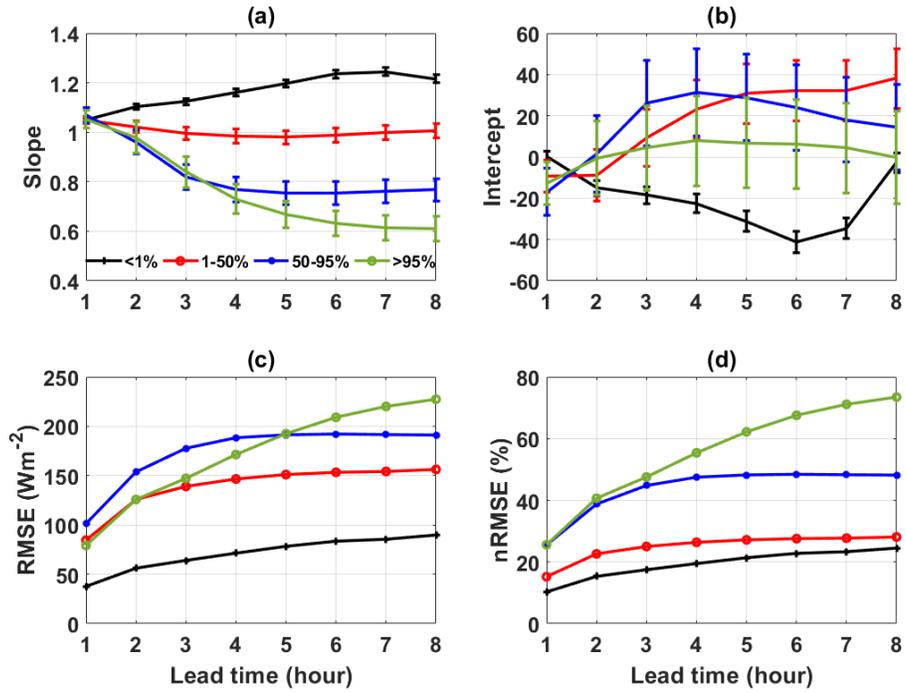
524



525

526 Figure 3. The LSTM performance for hourly averaged GHI prediction with different lead time
 527 from one to eight hours. (a) regression slope, (b) intercept, (c) RMSE, and (d) nRMSE. Results
 528 from different seasons are shown in different colors. The performances for the entire year are
 529 shown in black curves.

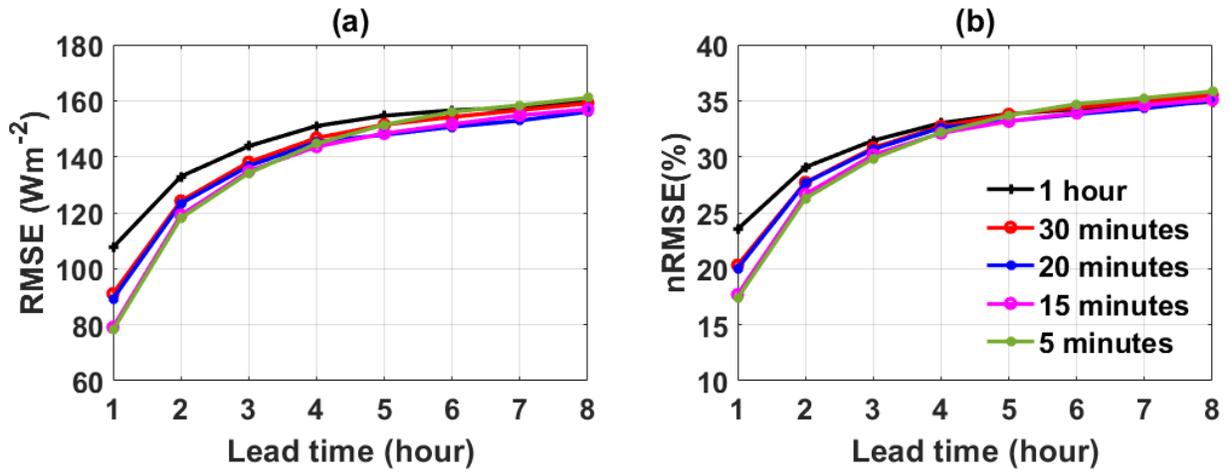
530



531

532 Figure 4. Performance of the LSTM forecast composite with respect to four groups of cloud
 533 fraction at forecast time zero: cloud fraction <1 % in red, 1-50 % in red, 50-95 % in blue and >95
 534 % in green. Four panels are the regression slope (a), regression intercept (b), RMSE (c) and
 535 nRMSE (d).

536

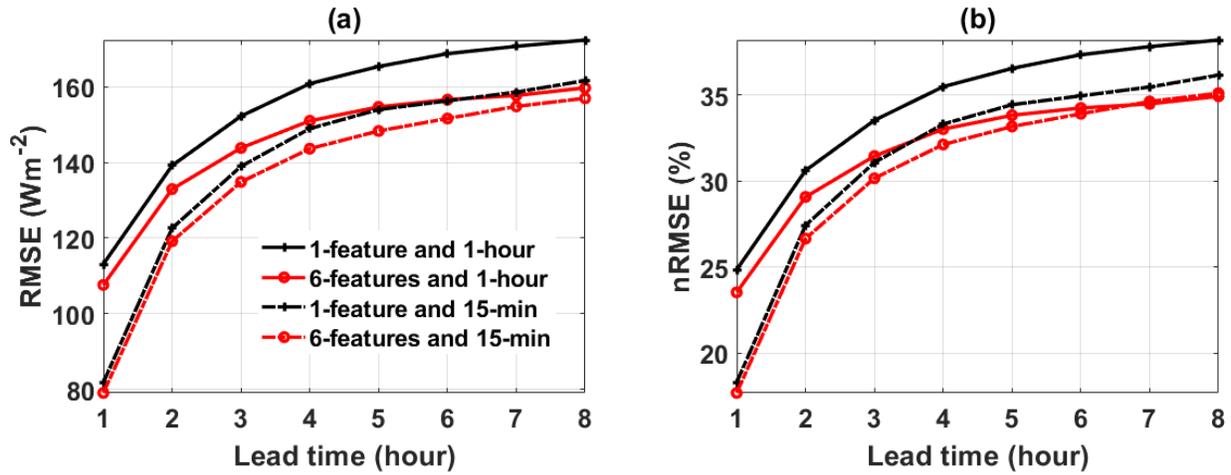


537

538 Figure 5. Sensitivity of RMSE (a) and nRMSE (b) to the temporal resolution of the inputs. The

539 same LSTM configuration used for Figures 2-4 is used here.

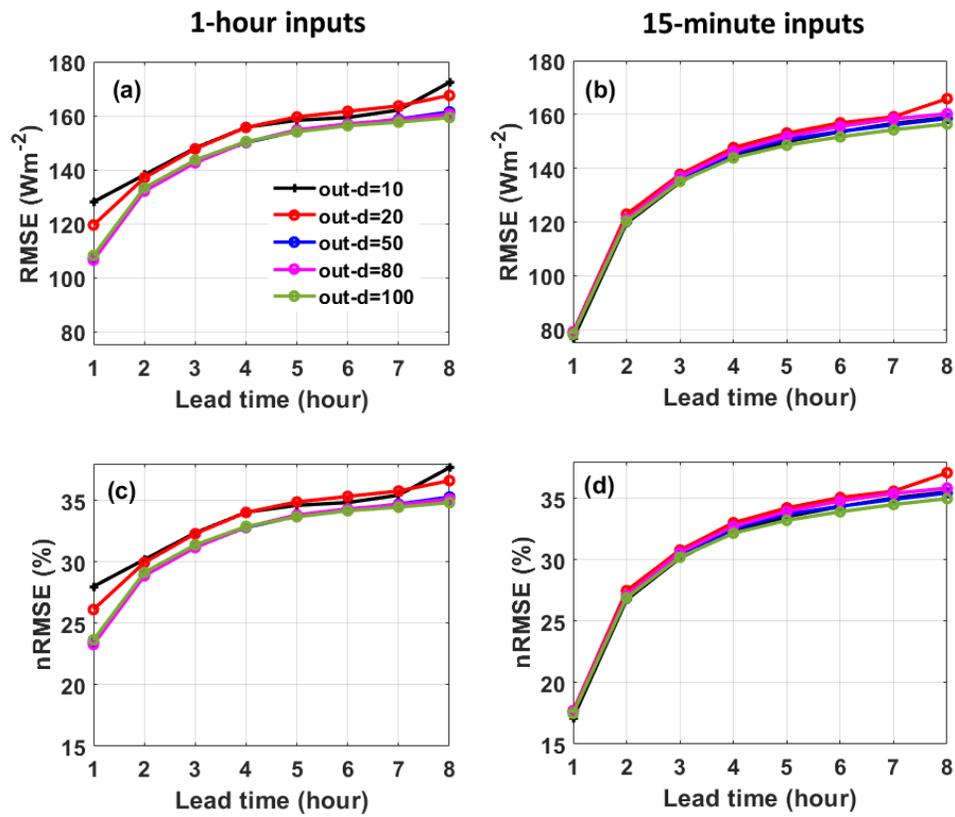
540



541

542 Figure 6. Sensitivity of RMSE (a) and nRMSE (b) to number of features used as the inputs. Results
 543 predicted using 1-hour inputs and 15-minute inputs are in solid and dashed lines, respectively. One
 544 feature is to use GHI only. Six features include time, GHI, total cloud fraction, solar irradiance at
 545 TOA, relative humidity, and air temperature. The same LSTM configuration used for Figures 2-4
 546 is used here.

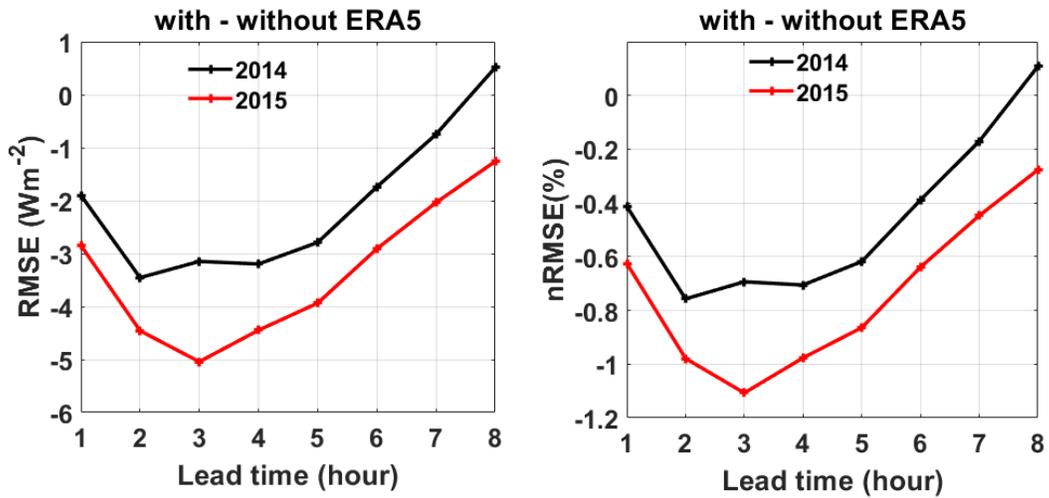
547



548

549 Figure 7. Sensitivity of RMSE and nRMSE to the number of LSTM output (out_d). Left column:
 550 1-hour inputs are used. Right column: 15-minute inputs are used. Upper row is for RMSE and
 551 lower row is for nRMSE. Results from different out_ds are shown in different colors, as labeled
 552 in (a) for all the panels.

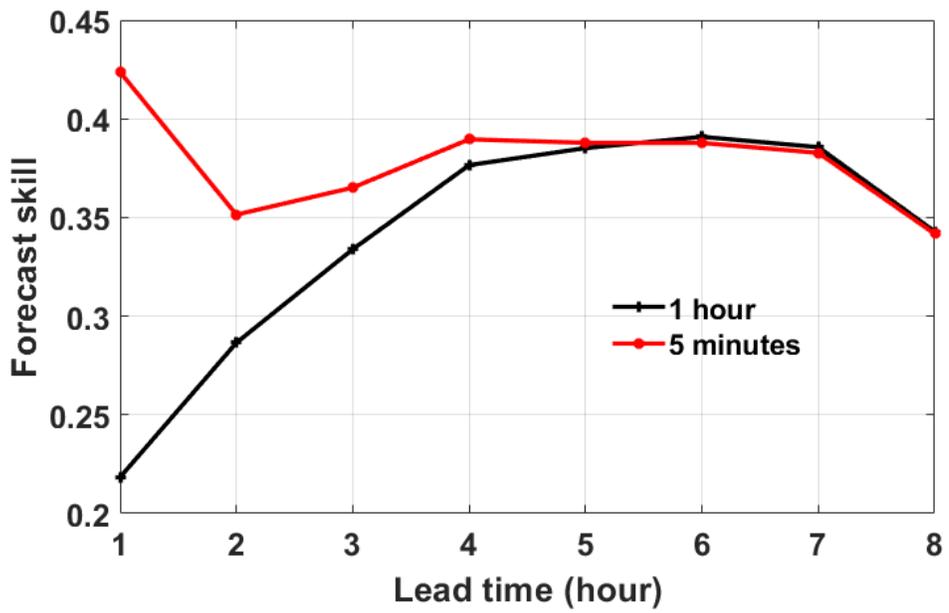
553



554

555 Figure 8. Difference in RMSE and nRMSE between with and without hourly ERA5 reanalysis
 556 used. Hourly inputs are used, other configurations are the same as in Figures 2-4. For black line,
 557 data in 2010-2013 are used to predict GHIs in 2014. For red line, data in 2011-2014 are used to
 558 predict GHIs in 2015.

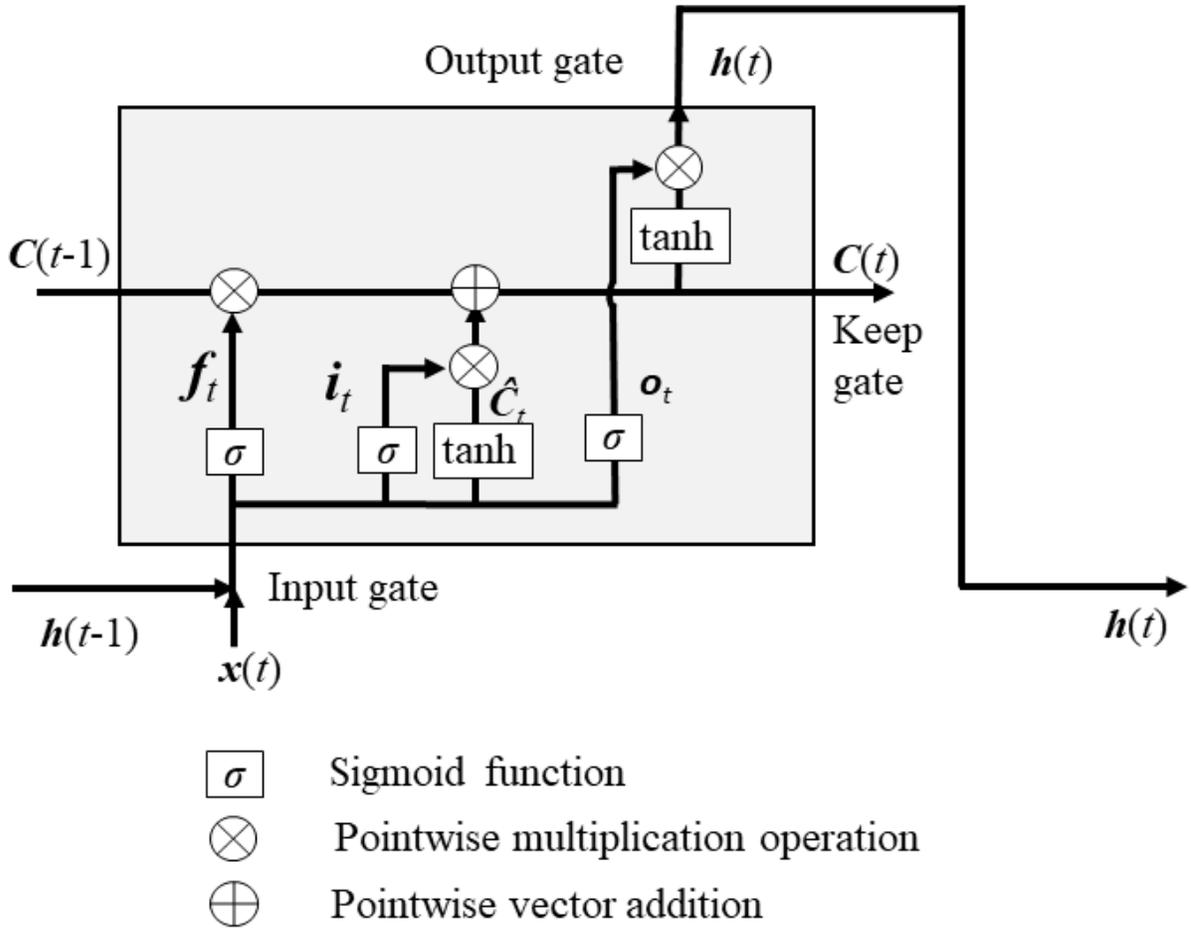
559



560

561 Figure 9. Forecast skill in 2014 compared to smart persistence model.

562



563

564 Figure A.1 Sketch of the structure of LSTM neural network cell based on illustration from Olah
 565 (2015) and Sun et al. (2017).

566