

Verification of deterministic solar forecasts¹

Dazhi Yang^{a,*}, Stefano Alessandrini^b, Javier Antonanzas^c, Fernando Antonanzas-Torres^c, Viorel Badescu^d, Hans Georg Beyer^e, Robert Blaga^f, John Boland^g, Jamie M. Bright^h, Carlos F.M. Coimbraⁱ, Mathieu David^j, Âzeddine Frimane^k, Christian A. Gueymard^l, Tao Hong^m, Merlinde J. Kayⁿ, Sven Killinger^o, Jan Kleissl^l, Philippe Lauret^l, Elke Lorenz^o, Dennis van der Meer^p, Marius Paulescu^f, Richard Perez^q, Oscar Perpiñán Lamigueiro^f, Ian Marius Peters^s, Gordon Reikard^t, David Renné^u, Yves-Marie Saint-Drenan^v, Yong Shuai^w, Ruben Urraca^c, Hadrien Verbois^h, Frank Vignola^x, Cyril Voyant^l, Jie Zhang^y

^aSingapore Institute of Manufacturing Technology, Agency for Science, Technology and Research, Singapore

^bResearch Application Laboratory, National Center for Atmospheric Research, Boulder, CO, United States

^cDepartment of Mechanical Engineering, University of La Rioja, Logrono, Spain

^dCandida Oancea Institute, Polytechnic University of Bucharest, Bucharest, Romania

^eFaculty of Science and Technology, University of the Faroe Islands, Tórshavn, Faroe Islands

^fFaculty of Physics, West University of Timisoara, Timisoara, Romania

^gCentre for Industrial and Applied Mathematics, University of South Australia, Mawson Lakes, SA, Australia

^hSolar Energy Research Institute of Singapore, National University of Singapore, Singapore

ⁱDepartment of Mechanical and Aerospace Engineering, University of California, San Diego, CA, USA

^jPIMENT Laboratory, University of La Reunion, Reunion, France

^kFaculty of Science, Ibn Tofail University, Kenitra, Morocco

^lSolar Consulting Services, Colebrook, NH, USA

^mDepartment of Systems Engineering and Engineering Management, University of North Carolina at Charlotte, Charlotte, NC, USA

ⁿSchool of Photovoltaic and Renewable Energy Engineering, University of South Wales, Sydney, NSW, Australia

^oFraunhofer Institute for Solar Energy Systems ISE, Freiburg, Germany

^pDepartment of Engineering Sciences, Uppsala University, Uppsala, Sweden

^qAtmospheric Sciences Research Center, University at Albany, SUNY, Albany, NY, USA

^rSchool of Engineering and Industrial Design, Polytechnic University of Madrid, Madrid, Spain

^sDepartment of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

^tStatistics Department, U.S. Cellular, Chicago, IL, USA

^uDave Renné Renewables, Boulder, CO, USA

^vMINES ParisTech, PSL Research University, Sophia Antipolis, France

^wSchool of Energy Science and Engineering, Harbin Institute of Technology, Harbin, Heilongjiang, China

^xMaterial Science Institute, University of Oregon, Eugene, OR, USA

^yDepartment of Mechanical Engineering, University of Texas at Dallas, Richardson, TX, USA

Abstract

The field of energy forecasting has attracted many researchers from different fields (e.g., meteorology, data sciences, mechanical or electrical engineering) over the last decade. Solar forecasting is a fast-growing subdomain of energy forecasting. Despite several previous attempts, the methods and measures used for verification of deterministic (also known as single-valued or point) solar forecasts are still far from being standardized, making forecast analysis and comparison difficult.

To analyze and compare solar forecasts, the well-established Murphy–Winkler framework for distribution-oriented forecast verification is recommended as a standard practice. This framework examines aspects of forecast quality, such as reliability, resolution, association, or discrimination, and analyzes the joint distribution of forecasts and observations, which contains all time-independent information relevant to verification. To verify forecasts, one can use any graphical display or mathematical/statistical measure to provide insights and summarize the aspects of forecast quality. The majority of graphical methods and accuracy measures known to solar forecasters are specific methods under this general framework.

Additionally, measuring the overall skillfulness of forecasters is also of general interest. The use of the root mean square error (RMSE) skill score based on the optimal convex combination of climatology and persistence methods is highly recommended. By standardizing the accuracy measure and reference forecasting method, the RMSE skill score allows—with appropriate caveats—comparison of forecasts made using different models, across different locations and time periods.

¹All co-authors are listed in alphabetical order.

46 *Keywords:* Solar forecasting, Measure-oriented forecast verification, Distribution-oriented forecast verification, Skill
47 score, Combination of climatology and persistence

48 1. Introduction

49 The power grids, which transmit and distribute electricity to end users, are being monitored and controlled by sys-
50 tem operators at all times to ensure reliable power delivery. Considering that solar and other renewable energy sources
51 are inherently variable, and that utility-scale energy storage is not economically viable globally yet, operational ex-
52 cellence of the power grids can benefit from accurate solar forecasts.² Consequently, reliable and well-characterized
53 solar forecasting tools and methodologies are becoming essential, and are considered of high value (Martinez-Anido
54 et al., 2016; Huang and Thatcher, 2017; Antonanzas et al., 2017; Klingler and Teichtmann, 2017).

55 Surface shortwave radiation is unavailable at night and, during daytime, fluctuates as a function of the position of
56 the Sun, cloud cover, aerosols, and other weather variables. Solar forecasts are used by utilities for various reasons:
57 switching energy sources, planning backup generators, calculating reserves, and energy trading. The time horizons
58 covered by modern solar forecasting typically range from a few seconds to a few days. Over the last decade, the
59 literature on this topic has bloomed. A wide spectrum of methods, either physics-based (e.g., sky or shadow imagery,
60 remote sensing, or numerical weather prediction), data-driven (e.g., time series, spatio-temporal statistics, or machine
61 learning), or a combination of both (e.g. hybrid models), have been proposed (see Blaga et al., 2019; Yang et al.,
62 2018; van der Meer et al., 2018; Voyant et al., 2017; Antonanzas et al., 2016; Ren et al., 2015; Inman et al., 2013,
63 for reviews). Furthermore, the existing studies span a range of time intervals and locations, with contrasting weather
64 conditions. Because of these differences, the field would benefit from having a general verification framework for
65 forecast analysis, as well as for the standardization of accuracy measures or metrics³ for forecast comparison.

66 This article has three missions. The first is to introduce the distribution-oriented forecast verification framework
67 to the solar forecasting community. The idea of using distributions—in particular the joint distribution of forecasts
68 and observations—originates in the work of Murphy and Winkler (1987). A joint distribution contains all time-
69 independent information relevant to verification. As such, it offers a more detailed view than the traditional measure-
70 oriented approach in terms of forecast analysis. The second mission is to recommend an accuracy measure that
71 should be universally reported in deterministic (also known as single-valued or point) solar forecasting studies—the
72 root mean square error (RMSE) skill score⁴ based on the optimal convex combination of climatology and persistence.
73 Since the skill score is able to reflect the inherent difficulties in different forecasting situations, it allows forecast
74 comparisons based on relative improvements, rather than on an absolute error size. The third mission is to look into a
75 series of practical issues in terms of forecast verification, such as data processing or implementation of the reference
76 forecasting methods, with the goal of helping users better understand the relative strengths and weaknesses of various
77 forecasting models in a uniform manner.

78 Even though the present authors represent a broad range of active researchers in the solar forecasting community,
79 there will always be difficulties in gaining universal consensus on the appropriate measures and methods to be used
80 in general, or even more so in various specific cases. Nevertheless, the authors hope that the forecast verification

*Corresponding author. Tel.: +65 9159 0888.

Email address: yangdazhi.nus@gmail.com (Dazhi Yang)

²Generally, the phrase “solar forecasting” refers to both solar power forecasting and solar irradiance forecasting. For the latter, forecasters are interested in two irradiance quantities, namely, global horizontal irradiance (GHI) and direct normal irradiance (DNI). Methodologically, there is not much difference in treating GHI and DNI. However, DNI forecasts are usually less accurate due to the larger variability in DNI than in GHI.

³“Measures” and “metrics” are distinct concepts in *measure theory*. A measure μ on a set X is a mapping $\mu : \mathcal{A} \rightarrow [0, \infty]$ defined on a σ -algebra \mathcal{A} that satisfies non-negativity, null empty set, and σ -additivity, that is $\mu(A) \geq 0 \forall A \in \mathcal{A}$, $\mu(\emptyset) = 0$, and $\mu(\sqcup_{j \in \mathbb{N}} A_j) = \sum_{j \in \mathbb{N}} \mu(A_j)$, where symbol \sqcup denotes disjoint union (Schilling, 2017). On the other hand, a metric is a distance measure $d : X \times X \rightarrow [0, \infty]$ that satisfies definiteness, symmetry, and triangle inequality, that is $d(x, y) = 0$ iff $x = y$, $d(x, y) = d(y, x)$, and $d(x, y) \leq d(x, z) + d(z, y)$, $\forall x, y, z \in X$. (Schilling, 2017). Nonetheless, moving out from measure theory, the two terms are often used interchangeably, e.g., “accuracy measure” and “error metrics” use the words “measure” and “metric” in their everyday sense. To most forecasters, especially forecast practitioners, they both refer to functions of forecast errors, such as mean bias error (MBE), mean absolute error (MAE), or root mean square error (RMSE).

⁴Skill score is also known as *forecast skill*. It is a class of accuracy measures that gauge the relative improvement of a method over a reference method, see Eqs. (3) and (4) below.

81 procedure proposed here can lead to a greater interpretability of results, and even in direct—“apples to apples”—
82 comparisons of techniques. An ultimate goal is to establish the best practices that can be upgraded and refined as the
83 body of knowledge and experience grows.

84 The organization of this study is as follows. The forecast verification problem and the perceived difficulties are
85 elaborated in Section 2. Distribution-oriented forecast verification methods are discussed and exemplified in Section 3.
86 The recommended accuracy measure is justified in Section 4, alongside with some discussions on practical concerns.
87 Section 5 concludes with a series of recommendations.

88 2. Problem description

89 Solar forecasting is a term applied to any predictive form of estimating the solar energy resource ahead of time.
90 With a fast-growing global portfolio of solar energy installations using various technologies, the need for solar fore-
91 casting to facilitate improved operations and electricity market compatibility is paramount. A rapidly expanding sci-
92 entific community in the subdomain of energy forecasting has contributed numerous methodologies and approaches
93 towards solar forecasting (Hong et al., 2016). Accuracy is a major goal of most, if not all, forecasters. The variability
94 in solar irradiance intrinsically governs predictability (Pedro and Coimbra, 2015). Therefore, it is particularly interest-
95 ing to compare forecasts generated by different models, using data from different locations,⁵ or different time periods
96 (Yang, 2019a).

97 Current methods of solar forecast verification are mostly limited to using measures as indicators of goodness of
98 forecasts. In other words, solar forecasters compare the performance of different models based on some error metrics,
99 and can then draw conclusions. Under this type of verification procedure, any conclusion is ambiguous in at least two
100 ways: (1) it is unclear what the forecast objective is, and (2) it is unclear how the model of interest performs against
101 other models that are not included in the study. These problems are described in sections 2.1 and 2.2, respectively.

102 2.1. What is a good forecast?

103 The word “objective” refers to goals given to a forecaster prior to verification. It is natural to think of the objective
104 as “small RMSE,” “high skill score,” or “high economic value.” Nonetheless, these objectives often lack generality,
105 and can even be conflicting at times. In that, one may end up collecting a large, and possibly redundant, set of error
106 metrics. This is exemplified by the work of Zhang et al. (2015), in which a suite of 17 metrics was assembled based on
107 a lengthy discussion process that involved stakeholders from both the meteorological and power systems communities.
108 In other cases, new metrics are proposed to meet a specific objective. This is exemplified by the work of Vallance
109 et al. (2017), in which the ability to forecast ramps in irradiance transients is gauged by two new metrics.

110 Assembling or introducing new members to a pool of error metrics is meaningful to the field of solar forecasting.
111 By presenting a wide spectrum of error metrics, forecasters are able to choose freely among the metrics that can “best”
112 highlight the strengths of their results. There are many studies that propose, contrast, and recommend error metrics
113 to forecasters (e.g., Vallance et al., 2017; Zhang et al., 2015; Hoff et al., 2013; Beyer et al., 2009). However, despite
114 the well-argued discussions, these works can rarely change another forecaster’s sentiment towards some specific
115 metrics, if they are perceived as having important advantages or disadvantages. Hence, for each argument that favors
116 a metric, one may find a counter-argument against it (see Chai and Draxler, 2014; Willmott and Matsuura, 2005).⁶
117 Furthermore, since there are countless publications that discuss and conclude that one metric is better than the other,
118 it is not difficult to cite those articles that support any choice the author wishes to make (Chai and Draxler, 2014). The
119 obvious consequence is a field with diverse usage of error metrics. Nonetheless, this is not unique to the emerging
120 field of solar forecasting. Historically, the lack of unified forecast verification procedure has been discussed by many

⁵Verification of forecasts, particularly those made by a numerical weather prediction (NWP) model, can be carried out spatially (see Gilleland et al., 2010). Spatial averaging and spatial scale have a strong impact on forecast accuracy (Lorenz et al., 2016). This work, however, is constrained to the verification of forecasts at point locations.

⁶During the initial stage of this study, the original idea was to propose a specific suite of metrics to the community. However, soon it became obvious that it is impossible to make everyone agree. No consensus can be established on things such as whether MAE or RSME should be favored, whether normalized metrics should be used, or with which quantity the RMSE should be normalized: the mean, maximum, 1000 W/m², or square root of second moment.

Table 1: Simulated clear-sky index (κ , dimensionless) and clear-sky GHI (c , in W/m^2) data. The κ simulation follows Eqs. (1) and (2).

Day 1			Day 2			Day 3			Day 4			Day 5		
Index	κ	c	Index	κ	c	Index	κ	c	Index	κ	c	Index	κ	c
1	1.000	169	12	0.987	170	23	0.824	176	34	0.914	182	45	0.691	168
2	0.941	389	13	0.928	389	24	0.749	397	35	0.993	403	46	0.954	380
3	0.997	587	14	0.987	589	25	0.577	597	36	0.955	601	47	0.683	576
4	0.815	739	15	0.919	743	26	1.000	752	37	0.972	753	48	0.989	728
5	0.871	832	16	0.730	836	27	0.561	846	38	0.770	845	49	0.999	820
6	0.999	859	17	0.393	864	28	0.968	874	39	0.912	871	50	0.998	845
7	0.999	817	18	0.873	822	29	0.972	832	40	0.973	827	51	0.895	800
8	0.963	710	19	1.000	716	30	0.952	724	41	0.995	719	52	0.938	687
9	0.995	546	20	0.925	552	31	0.993	558	42	0.989	553	53	0.648	516
10	0.710	340	21	0.984	346	32	0.960	351	43	0.714	345	54	0.960	308
11	0.998	121	22	0.671	124	33	0.992	130	44	0.918	125	55	0.871	105




Figure 1: A window of 55 simulated hourly GHI and clear-sky GHI data points (with zenith angle $< 85^\circ$). Forecasts generated by three forecasters over the same window are overlaid. The *novice* uses 1-step-ahead clear-sky persistence model, the *optimist* always uses 0.95 times of clear-sky GHI as forecasts, and the *statistician* uses the true conditional mean as forecasts.

121 experts from other relatively mature fields (e.g., [Murphy and Winkler, 1987](#); [Armstrong, 2001](#); [Fildes et al., 2008](#)),
 122 but nothing seems to have changed ([Gneiting, 2011](#)).

At this stage, it is essential to ask the question: “what is a good forecast?” It is known, *a priori*, that different metrics favor different forecasts. To put this issue in perspective, a simulation study is presented. Suppose diurnal variation of the hourly clear-sky index, i.e., the ratio between the global horizontal irradiance (GHI) and clear-sky GHI, at an arbitrary location follows:

$$\kappa_t = 1 - z_t^2, \quad (1)$$

where κ_t denotes the clear-sky index at time t , $z_t \sim \mathcal{N}(0, \sigma_t^2)$, and σ_t^2 follows:

$$\sigma_t^2 = 0.15z_{t-1}^2 + 0.3\sigma_{t-1}^2 + 0.07. \quad (2)$$

123 With initial values $z_0 = 0$ and $\sigma_0^2 = 0.01$, the κ time series is simulated for 55 daylight hours (herein defined to be
 124 data points with a zenith angle $< 85^\circ$), or 5 days. The simulated data points are tabulated and plotted in [Table 1](#) and
 125 [Fig. 1](#), together with the corresponding clear-sky GHI (c) values. The McClear model ([Lefèvre et al., 2013](#)) is used to
 126 estimate c .

127 Based on the simulated time series, three forecasters are asked to generate forecasts. The *novice* has no skill to
 128 offer, and thus issues 1-step-ahead persistence forecasts on κ , i.e., $\phi_t = \kappa_{t-1}$, where symbol “ ϕ ” denotes the forecast
 129 clear-sky index. The *optimist* knows it is sunny in that location, and always uses $\phi_t = 0.95$. The *statistician* has
 130 knowledge about the inherent model, and thus issues the true conditional mean as forecasts, i.e., $\phi_t = 1 - \sigma_t^2$. These ϕ
 131 values are then converted to GHI forecasts with clear-sky GHI values at the forecast timestamps. The results in terms

132 of three error metrics, namely, mean bias error (MBE), mean absolute error (MAE), and RMSE⁷ are shown in Table 2.
 133 The results are inconclusive, because each forecaster is best in terms of a particular error metric.

Table 2: MBE, MAE, and RMSE, in W/m², of the three forecasters in the simulation study. Column-wise best results are in bold.

Forecaster	MBE	MAE	RMSE
Novice	-1.32	79.80	127.12
Optimist	33.45	53.96	100.51
Statistician	11.43	58.08	93.97

134 The result of the above simulation study contradicts the common belief that knowing the inherent (physical or
 135 statistical) process is the determining factor behind making good forecasts. This contradiction is attributed to how
 136 the goodness of forecasts is defined. To most solar forecasters, “good forecast” is implicitly equivalent to “small
 137 error.” However, the pitfalls of this definition become apparent whenever contradicting rankings of models materialize.
 138 In order to resolve such contradictions, solutions might be obtained from the field of meteorology, where forecast
 139 verification is well studied.

140 [Murphy \(1993\)](#) outlined three types of goodness that jointly define a good forecast:

- 141 1. *consistency*—correspondence between forecasts and judgments;
- 142 2. *quality*—correspondence between forecasts and observations; and
- 143 3. *value*—incremental benefits of forecasts to users.

144 2.1.1. Consistency

145 *Consistency* is quite an abstract concept: a forecast is consistent if it corresponds with the forecaster’s best judg-
 146 ment. [Murphy \(1993\)](#) argued that such a judgment must contain an element of uncertainty, because the forecaster’s
 147 knowledge on the forecasting task is necessarily incomplete. In probabilistic forecasting, consistency can be ensured
 148 by adopting *strictly proper scoring rules* ([Gneiting and Raftery, 2007](#)). With that, forecasters are rewarded with the
 149 best scores if and only if their forecasts correspond with their judgment ([Murphy and Winkler, 1971](#)). The Brier score
 150 and continuous ranked probability score (CRPS), both of which frequently used in probabilistic solar forecasting, are
 151 both strictly proper ([Gneiting and Raftery, 2007](#)).

152 On the other hand, in deterministic forecasting, forecasters have to translate their probabilistic judgment through
 153 a statistical functional,⁸ $T(F)$, which summarizes the forecast distribution, F . The reader is referred to [Gneiting](#)
 154 (2011) for the formal definition. Informally, the scoring function S is consistent if $\mathbb{E}[S(f, x)] \leq \mathbb{E}[S(g, x)]$, for all
 155 $f \in T(F)$, where f is an evaluation of the functional, g is any forecast, and x is a future observation. This definition
 156 implies that S is consistent if and only if any $f \in T(F)$ is an optimal forecast under S . For example, if the mean
 157 value of a forecaster’s judgmental probability distribution is of interest, then RMSE is a consistent accuracy measure,
 158 because RMSE is minimized by forecasting the mean of the predictive distribution. In the above simulation study,
 159 the *statistician* provided the optimal forecasts under RMSE. The *optimist*, although winning the competition with
 160 respect to MAE, did not provide optimal forecasts under MAE (MAE is minimized by forecasting the median of the
 161 predictive distribution).

162 The underlying assumption of using consistency as a measure of goodness of forecasts is that the forecasters
 163 receive a *directive* in the form of a statistical functional to transform their probabilistic judgment into a deterministic
 164 forecast. For instance, the directive could be “forecast the mean of your probabilistic judgment.” Only then, a scoring
 165 rule can be identified as consistent if it is optimized by the chosen directive. However, [Jolliffe \(2008\)](#) noted that

⁷By surveying 1000 recent forecasting papers, [Yang et al. \(2018\)](#) found that there are about 20 commonly used metrics in solar forecasting, with MBE, MAE, and RMSE being the most popular ones. They can thus be considered most typical, which is why they are used in the simulation study. If the forecast and observed GHI at time t are denoted using f_t and x_t , respectively, then $\text{MBE} = \sum_{t=1}^{55} (f_t - x_t)$, $\text{MAE} = \sum_{t=1}^{55} |f_t - x_t|$, and $\text{RMSE} = \sqrt{\sum_{t=1}^{55} (f_t - x_t)^2}$

⁸Any function of a probabilistic distribution is called a statistical functional. Examples of functionals include mean, median, or variance. Generally, it is written as $T(F)$, where F is a distribution.

166 the very definition of “consistency” is circular: a forecaster can elect to start by choosing a scoring rule. Once the
167 forecasts are made by optimizing the scoring rule, the consistent directive naturally follows.

168 Consistency implies a theoretical guideline for choosing the most appropriate accuracy measure during forecast-
169 casting. For most statistical and machine-learning models, the model parameter or weights are estimated or fitted
170 according to some cost function. In this regard, a consistent error measure should be used during verification. For
171 instance, the ordinary least squares regression minimizes the sum of squared errors, hence, RMSE is an appropriate
172 metric to report. Nonetheless, such guideline might not favor forecast comparison in practice, since it would further
173 divide the field. Further studies are necessary.

174 2.1.2. Quality

175 *Quality* is a familiar concept to all solar forecasters, as it refers to the correspondence between forecasts and
176 observations. For example, MAE and RMSE are both measures that assess the overall accuracy of forecasts. Accuracy
177 is an aspect of forecast quality. It can be interpreted through quantitative measures. Besides accuracy, other aspects
178 of forecast quality known to solar forecasters, such as bias, association, skill, or uncertainty, can be assessed through
179 MBE, correlation, skill score, or variance. In forecast verification, the traditional way of comparing measures, may it
180 be positively oriented (the larger, the better, such as skill score), negatively oriented (the smaller, the better, such as
181 RMSE), or center oriented (the closer to a center value, the better, such as MBE), is known as the *measure-oriented*
182 *approach*.

183 As mentioned earlier, one disadvantage of using the measure-oriented approach is the subjectivity in choosing
184 measures. Since selecting which measures to report is essentially a decision that is internal to a forecaster, the reasons
185 behind that selection are by default unknown to anyone who is observing the forecast verification procedure from
186 an external view point. In academia, forecasters are authors, whereas observers are editors, reviewers, and readers.
187 If the *optimist* in the simulation study only reports MAE in an article, the observers will not be able to fully realize
188 the underlying pitfalls of those forecasts, but will have no other choice than to accept their results. The simulation
189 study above might over-simplify the state-of-the-art solar forecasting scenarios in comparison with actual forecasting
190 models. Because those are typically much more complex, it would be even more difficult to interpret their results
191 through only a few measures.

192 A related question is, if two forecasting methods yield the same MBE, RMSE, or skill score, are they equally
193 good? The obvious answer is “no.” Measures only provide an *overall* assessment of forecast quality. Since error
194 metrics are often computed based on a collection of samples (e.g., rolling hourly forecasts made over a year), this
195 gives infinite ways to result in the same error-metric value. The reader is referred to Fig. 1 in [Vallance et al. \(2017\)](#)
196 for an example on how two sets of drastically different forecasts can lead to the same RMSE. One solution frequently
197 being used in the solar forecasting literature is to report the regime-dependent error metrics, i.e., to differentiate the
198 errors by classes of prevailing situations. For instance, one can separately report errors for overcast-, clear-, and all-
199 sky conditions. Alternatively, one can also report errors for different times of day, different times of year, or different
200 day types. However, the dimensionality of forecast verification scales with the number of classes, e.g., an RMSE
201 table will become three, if three sky conditions are analyzed separately, or ten, if ten day types are defined. The error
202 contingency table often gets out of control quickly. What has just been discussed is known as *forecast analysis*, which
203 is generally defined as the procedure selected to understand the *composition* of the overall quality.⁹

204 Since both assessment and analysis of forecast quality are driven by the information embedded in the forecast-
205 observation pairs, it is useful to define the total amount of information available to a forecaster during verification.
206 By defining the entirety of information, a forecaster is no longer limited by the set of summary statistics. Stated
207 differently, if the temporal sequence of forecast-observation pairs is not of interest, the joint distribution of forecast
208 and observation can be used to study the skillfulness of the forecasts, since it contains *all* time-independent information
209 relevant to forecast verification. This *distribution-oriented approach* to forecast verification was proposed by [Murphy](#)
210 [and Winkler \(1987\)](#). It has gained high popularity in the field of meteorological forecasting, but is less known by solar
211 forecasters.

212 This particular framework needs to be discussed because it provides an alternative view to the traditional measure-
213 oriented approach. It offers high flexibility in terms of accessing the information. In fact, the majority of graphical

⁹A related issue is how to decompose these overall quality metrics. Some options are discussed in Section 3.

214 methods (e.g., Taylor diagram, target diagram, or error heat map) and accuracy measures (e.g., MBE, RMSE, or
215 Kolmogorov–Smirnov test integral) known to solar forecasters are specific methods under this general framework.
216 More importantly, the Murphy–Winkler framework is augmented by Bayes’ theorem, in that the joint distribution
217 can be written equivalently as the product of marginal and conditional distributions, making the embedded infor-
218 mation more accessible. Last but not least, the distribution-oriented approach establishes communication between
219 forecast quality and accuracy measure. Aside from those aspects of forecast quality mentioned earlier, other aspects
220 such as reliability, resolution, or discrimination can easily be defined and quantified. Whereas more details on the
221 Murphy–Winkler framework are provided in Section 3, with a case study, it is noted that the framework is essential to
222 understanding the goodness of forecasts.

223 2.1.3. Value

224 *Value* relates to the benefits realized, or cost incurred, by individuals or organizations who use the forecasts during
225 their decision making. Murphy (1993) pointed out that the forecasts by themselves possess no intrinsic value, because
226 they only acquire value through influencing the decisions made by their users. Most often, the value of solar forecasts
227 is translated into and measured in monetary units.¹⁰ For instance, by reducing the RMSE of the forecasts by x W/m²,
228 the owner of a photovoltaic (PV) system with energy storage may expect to gain an additional \$ y per year through
229 optimizing the feed-in strategy of the system. In a detailed case study, Law et al. (2016) discussed the benefits of
230 improvements in irradiance forecasting for a concentrated solar thermal power plant in this context. An alternative
231 view was given by Antonanzas et al. (2017), where they compared the profit from different forecasting methods with
232 respect to that from perfect forecasts.

233 Naturally, such benefits or costs depend on the characteristics of a particular decision-making problem. Thus, the
234 third type of goodness is not under the control of forecasters, but is determined and appreciated by decision makers.
235 Furthermore, this goodness of forecast is non-transferable by default. That is, one cannot simply scale the value
236 realized by others, using the characteristics of the problem at hand. Because of the different courses of action and
237 payoff structures available to different decision makers, there is little reason to assume an *ex post* value would apply
238 in an *ex ante* study. A good forecasting strategy that creates high value to some users might not be appreciated by
239 others.

240 In parallel, it is believed that for a fixed and well-defined decision-making problem, the mapping between quality
241 and value is *monotone*. In other words, higher forecast quality expectedly corresponds to higher value. In principle,
242 this gives the forecaster the necessary motivation to provide the best possible (and hopefully “optimal”) forecasts.

243 To give a perspective on what a “well-defined decision-making problem” can actually be, the case of the Australian
244 National Electricity Market (NEM) is considered. In NEM, conventional generators submit bids every five minutes
245 to the Australian Energy Market Operator. The latter tries to see how far up the bid stacks they have to proceed
246 to meet their forecast net load (regional load forecast minus forecast of domestic and commercial PV generation).
247 If the conventional generators miss their promised amount by more than a given tolerance, either above or below,
248 they are penalized. There is presently a dramatic expansion in solar farm construction and these installations are
249 price takers, i.e., not involved in making the spot market price. Hence, the solar plants would not be fined for poor
250 forecasts, but could be curtailed whenever necessary. Under this regime, the decision-making problem might not be
251 well defined, because the plant owners could always use the highest possible power generation as forecasts, since there
252 is no monetary penalty imposed on over-forecasts. To level the playing field in a new regime, the solar plants could be
253 penalized too, if they do not meet their forecasts. Given the new payoff structure, the cost of over-forecasts can be set
254 to be equivalent to the cost of running spinning reserve to fulfill the difference between the forecast and the generated
255 solar electricity. Similarly, the cost of under-forecasts would be the lost revenue that would have been generated if
256 that extra potential energy were sold at the prevailing spot price at the time. In this case, the decision-making problem
257 is well defined *if* the two costs are the same, then encouraging the forecasters to submit their optimized forecasts
258 truthfully.

¹⁰Other measures of value could be in terms of stability of the grid, customer satisfaction, etc. But how the value is measured does not affect the discussion here.

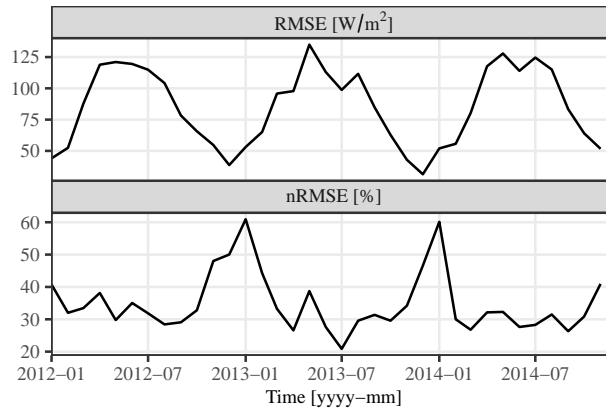


Figure 2: Monthly averaged day-ahead GHI forecast errors—RMSE [W/m^2] and nRMSE [%] normalized by monthly mean irradiance values—in Hamburg, Germany (53.63°N , 9.99°E), produced by the Integrated Forecasting System (IFS), an NWP system developed by the European Centre for Medium-Range Weather Forecasts (ECMWF). It is evident that winter months have small RMSE but high nRMSE, which might be confusing.

2.2. The skill score

The three types of goodness defined by [Murphy \(1993\)](#) provide a clear objective during forecast verification—while maintaining consistency, one should aim at maximizing quality. However, having a well-defined objective only helps the forecasters to analyze and thus make conclusions based on their own forecasting experiment. As the specialized literature expands rapidly, it is unrealistic to expand the scope of the experiment simultaneously, that is, to include all previously proposed methods as benchmarks. The obvious reason is that the data (information) available to one forecaster might not be available to others. Similarly, not all types of information available at one location or time is available at other places or times. Thus, to guarantee progress in the field, the community is forced to make comparison among different research works, based on a variety of reported measures of forecast quality.

2.2.1. A false sense of cross-scenario comparability

The variability of solar irradiance depends on climatic features, geographical location, timescale, and time period, among other factors. Even if the same forecast-generating strategy is employed, the hourly forecasts made for a location with predominant clear-sky conditions will have a significantly smaller RMSE as compared to 10-min forecasts made for a site with a tropical climate, where cloud formation is rapid and difficult to predict. Hence, if one wishes to compare forecast skills, some form of scaling (normalization) is needed for scale-dependent errors, such as MAE or RMSE. In a recent review paper, [Blaga et al. \(2019\)](#) used the normalized RMSE (nRMSE) as a basis for such comparisons.

The particular form of normalization considered in [Blaga et al. \(2019\)](#) is through the mean of the observations, i.e., nRMSE is computed by dividing RMSE (e.g., in irradiance unit) with the mean of the observations. Whereas the final conclusion—nRMSE reported in the solar forecasting literature is getting smaller over the years—is factual, the methodology (directly comparing nRMSE results from various publications) used by the authors can be misleading. “Researchers are getting better at forecasting solar” is *a priori* knowledge, and it is easy to find evidence supporting that. However, nRMSE gives a false sense of cross-scenario comparability, which cannot be used to justify that one forecaster has better skill than the other. The word “cross-scenario” refers to forecasts generated for data with different predictability.¹¹ One example is shown in Fig. 2, where RMSE and nRMSE lead to contradictory conclusions.

As mentioned earlier, forecast error is tightly linked to predictability, which is related to variability and uncertainty. Equivalently, it can be stated that conditions with higher variability and uncertainty are harder to forecast, thus leading to larger expected errors. In forecasting, variability and uncertainty are often quantified by step change and variance, respectively. Since the ultimate aim is to have a measure that quantifies forecast skill, its dependency on variability

¹¹If two sets of forecasts are generated for datasets with the same variability, e.g., two PV systems with different nominal power at the same location, nRMSE is a good normalization choice.

288 and uncertainty has to be minimized if not removed completely. It is now clear that mean-normalized nRMSE cannot
 289 be used to compare forecast skills in general, because the mean is related to neither variability nor uncertainty. The
 290 same arguments can be applied to range-normalized nRMSE, max-normalized nRMSE, capacity-normalized nRMSE,
 291 and similar variants of nMAE.

292 2.2.2. *On the propagation of normalized accuracy measures in solar forecasting*

293 Normalized accuracy measures are popular in solar forecasting (Blaga et al., 2019). This is in contrast to the
 294 field of meteorology, where normalized accuracy measures are rarely used. For instance, in the book by Jolliffe and
 295 Stephenson (2012), there is not a single sentence that discusses normalized accuracy measures. Similarly, no trace
 296 of normalized accuracy measures can be found in Hyndman and Koehler (2006), in which the accuracy measures are
 297 discussed in the context of general-purpose univariate time series forecasting. Hence, some *possible* explanations on
 298 why normalized accuracy measures is so popular in solar forecasting are given next.

299 The notion of normalization develops naturally when the forecast quantities are at different scales. On this point,
 300 the class of accuracy measures based on percentage errors, such as the mean absolute percentage error (MAPE),
 301 needs to be discussed. Measures based on percentage errors are not quite feasible in high-resolution solar forecasting
 302 since the irradiance and PV-generated power is near zero during early mornings and late afternoons, or when the
 303 clouds move in. Although the early morning and late afternoon cases can be trimmed with a zenith-angle filter before
 304 verification, a few missed forecasts on large irradiance swings during mid-day are enough to result in a very large
 305 MAPE. Therefore, to allow the forecast errors to be interpreted as a realistic percentage, the normalization is taken
 306 out of the summation, i.e., normalization is performed after aggregation. No such concerns exist when dealing with
 307 longer-term forecasts of global irradiance, e.g., on a daily basis.

308 Normalized accuracy measures are used in wind forecasting, a more developed sub-domain of energy forecasting
 309 (Hong et al., 2016). In an effort to standardize metric usage in wind forecasting, Madsen et al. (2005) noted that the
 310 purpose of using normalized accuracy measures is to produce results independent of wind farm sizes. In addition, the
 311 authors recommended normalization by the installed capacity or mean observation. At that time (2005), published
 312 studies on solar forecasting were rare. When the field of solar forecasting started to bloom in the early 2010s, such
 313 normalization was adopted by solar forecasters, to whom wind forecasting was the most relevant literature to follow.

314 Another compelling reason why normalized accuracy metrics are frequently used in solar (or wind) forecasting is
 315 that the end users (or “stakeholders”) are typically electrical engineers, business analysts or financial experts. These
 316 professionals are very familiar with percents, but not with a solar radiation unit such as W/m² or a power unit such
 317 as MW. From this standpoint, the use of normalization is essentially dictated by the necessity for the end users to
 318 understand and correctly use the forecast results. Nonetheless, grid operators almost never compare their forecasts
 319 across seasons, timescales, and let alone to forecasts of other grid operators. Therefore, in that context, the choice
 320 of normalized accuracy metrics is for convenience and internal communications, since a percentage metric is more
 321 accessible to a non-technical audience (decision-makers) than a MW metric. Since for grid operators the normalizing
 322 quantity, i.e. the denominator, is either constant (peak load) or similar (average load), the normalization does not
 323 affect the ranking of forecast accuracy. The grid operator preference for normalized accuracy metrics should *not* be
 324 construed as a decisive motivation to report only normalized error metrics in the academic community, because the
 325 global inter-comparability of forecast quality needs to be maintained.

326 2.2.3. *Problems with the skill score*

Since normalized accuracy measures cannot be used to compare forecasts made at different locations and timescales,
 an alternative has to be sought. In modern solar forecasting, one of the first attempts to formally address the problem
 of comparability was made in the early 2010s by Marquez and Coimbra in several publications (Marquez and Coim-
 bra, 2011, 2013; Coimbra et al., 2013).¹² In those studies, a well-known concept in meteorological forecasting called
skill score was introduced to the young field of solar forecasting. In the field of meteorology, the skill score, s , can be
 defined based on some measure of accuracy A , namely,

$$s = \frac{A_f - A_r}{A_p - A_r}, \quad (3)$$

¹²Another early discussion was given by Beyer et al. (2009).

where A_f , A_p , and A_r are the accuracy of the forecasts of interest, accuracy of the perfect forecasts,¹³ and accuracy of the reference forecasts, respectively (Murphy, 1988). For instance, s based on RMSE is

$$s = 1 - \frac{\text{RMSE}(f, x)}{\text{RMSE}(r, x)}, \quad (4)$$

where f , r , and x are forecasts of interest, reference forecasts, and observations, respectively. For N samples,

$$\text{RMSE}(f, x) = \sqrt{\frac{1}{N} \sum_{t=1}^N (f_t - x_t)^2}, \quad (5)$$

$$\text{RMSE}(r, x) = \sqrt{\frac{1}{N} \sum_{t=1}^N (r_t - x_t)^2}. \quad (6)$$

327 The skill score s is often written as a percentage, representing the percent improvement in accuracy of the forecasts
 328 over the forecasts produced by a reference method. If $s > 0$, the forecasts of interest have a smaller RMSE than
 329 that of the reference forecasts, otherwise, $s \leq 0$ indicates that the model of interest fails to outperform the reference
 330 forecasts. There are, however, two problems with using s to compare forecasts: (1) the choice of accuracy measure
 331 can be arbitrary, and (2) the choice of the reference forecasting method can be arbitrary.

332 The first problem can be understood with a simple example. The computation of s requires a measure of forecast
 333 accuracy, A , which is based on a scoring function. Depending on the choice of A , s can be quite different. For instance,
 334 suppose $\text{RMSE}(r, x) = 200 \text{ W/m}^2$ and $\text{RMSE}(f, x) = 100 \text{ W/m}^2$, then $s_{\text{RMSE}} = 0.5$. However, when the mean square
 335 error (MSE) is used, $s_{\text{MSE}} = 1 - (1 - s_{\text{RMSE}})^2$ is boosted to 0.75. Whereas the conversion between s_{RMSE} and s_{MSE} is
 336 straightforward, s calculated based on other metrics, such as MAE, would be different, and cannot be inferred from
 337 s_{RMSE} or s_{MSE} . Hence, there is no obvious solution to this but to recommend a consensus. At the moment, RMSE is
 338 the most common form of A in the literature (Blaga et al., 2019; Yang et al., 2018), and thus should be predominantly
 339 used in skill score computation. (This choice is discussed further in Section 4.1). Hereafter, the symbol s only denotes
 340 s_{RMSE} , unless otherwise stated.

341 One remedy to the second problem is to use a universally-accepted naïve reference model, so that s can be used—
 342 with appropriate caveats—to compare the accuracy of forecasts made across different locations or time periods. Skill
 343 score is built upon the notion that the forecast errors of a “no skill” reference method should sufficiently reflect the
 344 difficulty of the forecasting scenario. In business forecasting, random walk is often used as the reference, and the
 345 relative performance of the model of interest is gauged using the Theil’s U statistic, a concept similar to skill score
 346 (Makridakis et al., 2008). In meteorology, the so-called “climatology” is often used as the naïve reference (Jolliffe
 347 and Stephenson, 2012). In deterministic solar forecasting, one of the most popular naïve reference methods—for
 348 intra-hour and intra-day forecasts—is the clear-sky adjusted persistence, or simply, *clear-sky persistence*. Clear-sky
 349 persistence is conceptually no different from the seasonal naïve method described in Makridakis et al. (2008). That is,
 350 persistence forecasts are made based on the clear-sky index, and then adjusted back to irradiance using the clear-sky
 351 irradiance at each forecast timestamp. The clear-sky model used in clear-sky persistence can effectively describe two
 352 seasonal cycles (a yearly cycle and a daily cycle) in an irradiance time series. Some alternative seasonality adjustment
 353 approaches are discussed in Appendix A. As for day-ahead solar forecasting, the *24-h persistence*¹⁴ and the *daily-*
 354 *average climatology* are popular. In this paper, an additional naïve reference, namely, the optimal convex combination
 355 of climatology and persistence is introduced, which is guaranteed to outperform both individual reference methods.
 356 Such recommendation is addressed at length in Section 4.2.

¹³It should be noted that Eq (4) assumes the RMSE of the perfect forecast accuracy, A_p , is 0. However, in almost all statistical forecasting frameworks, the models would assume some unpredictable white noise, i.e., non-zero RMSE even if a model perfectly describes the data-generating process. Hence, the assumption here is that the $A_p \ll A_r$, so that it can be neglected.

¹⁴The term “24-h persistence” means the forecast at hour t on the day of interest is the same as the observation made at hour t on the previous day. If the typical operational forecast submission lead time is considered, the method rather needs to be referred to as “48-h persistence,” since the observations made two days prior to the operating day will be used as forecasts.

357 *2.3. The skill score defined by Marquez and Coimbra (2011)*

358 The skill score is not limited to the verification of deterministic forecasts of continuous random variable. It is
 359 also used in verification of deterministic forecasts of binary events (e.g, Gilbert skill score or Doolittle skill score),
 360 multi-category events (e.g., Gandin and Murphy score), and probabilistic forecast verification (e.g., Brier skill score
 361 or CRPS skill score), as described by Jolliffe and Stephenson (2012). While the reader is referred to Jolliffe and
 362 Stephenson (2012) for more details on the skill score concept, the version proposed by Marquez and Coimbra (2011)
 363 needs to be discussed. In deterministic solar forecasting, their version is one of the most notable alternatives to the
 364 skill score defined in Eq. (4).

Marquez and Coimbra (2011) proposed their skill score, denoted here as s^* , based on the concept of “uncertainty” (U) and “variability” (V):

$$s^* = 1 - \frac{U}{V}, \quad (7)$$

where

$$U = \sqrt{\frac{1}{N} \sum_{t=1}^N \left(\frac{f_t - x_t}{c_t} \right)^2}, \quad (8)$$

$$V = \sqrt{\frac{1}{N} \sum_{t=1}^N \left(\frac{x_t}{c_t} - \frac{x_{t-1}}{c_{t-1}} \right)^2}, \quad (9)$$

365 and f , x , and c are forecast, observation, and clear-sky expectation, respectively. Marquez and Coimbra noted that
 366 the ratio between U and V can be approximated by the ratio of the RMSE of the model of interest and the RMSE
 367 of clear-sky persistence, i.e., $s^* \approx 1 - \text{RMSE}(f, x) / \text{RMSE}(r, x) = s$. However, no detailed theoretical support was
 368 given in the different versions of that proposal (Marquez and Coimbra, 2011, 2013; Coimbra et al., 2013). Instead,
 369 the approximation was demonstrated empirically, using results from several time series models.

370 By comparing s^* to s defined in Eq. (4), it is not apparent why should the approximation between the two skill
 371 scores hold. In fact, s is the RMSE skill score of irradiance (or solar power) forecasts, whereas s^* is the RMSE skill
 372 score of clear-sky index forecasts. Stated differently, the two scores verify different forecast quantities—the former
 373 verifies irradiance forecasts, and the latter verifies clear-sky index forecasts. However, upon careful analysis, one can
 374 arrive at the conclusion: if the κ forecast error and κ reference forecast error are both independent of the clear-sky
 375 expectation, then $s^* = s$, see Appendix B. That said, such independence assumption is almost always violated to a
 376 certain degree, due to the imperfect clear-sky model, high sensitivity to three-dimensional effects for broken clouds
 377 situations, and the high sensitivity to errors in the atmospheric turbidity. Therefore, the clear-sky index forecast errors
 378 are often larger for low solar elevations (Sengupta et al., 2015). Consequently, one only observes $s^* \approx s$ in most cases.

379 *2.4. Section summary*

380 This section has discussed and put forward a few new concepts related to deterministic solar forecast verifica-
 381 tion. In general, there are two approaches for forecast verification, namely, the measure-oriented approach and the
 382 distribution-oriented approach. The goodness of forecasts contains three elements: (1) consistency, (2) quality, and
 383 (3) value. Particularly interesting is that the distribution-oriented approach can help forecasters assess the quality of
 384 their forecasts in a systematic way—by relating various aspects of forecast quality to different accuracy measures—so
 385 that they can be interpreted. This is discussed further in Section 3.

386 Whereas the distribution-oriented approach is primarily recommended for forecast analysis, i.e., to be used within
 387 a forecasting case study, the skill score is recommended for cross-work forecast comparison. It is important to note
 388 that the measure-oriented and distribution-oriented approaches are complementary and not substitutive. Skill score is
 389 computed based on the accuracy measure of a reference model that can sufficiently describe the difficulty (variability
 390 and uncertainty) of a forecast situation. It gauges the overall skillfulness of the forecaster. Considering the importance
 391 and economic consequences of inter-comparing a variety of forecasting models in practice, the RMSE skill score
 392 should be mandated and standardized. It should be noted that whenever climatology, persistence, or their combination

Table 3: Definition, interpretation, and quantification of Murphy–Winkler factorizations (Jolliffe and Stephenson, 2012; Murphy and Winkler, 1987).

distribution	definition	interpretation	(some specific methods of) quantification
$p(x f)$	(related to <i>calibration</i>)	A set of deterministic forecasts is perfectly calibrated if $\mathbb{E}(x f) = \int xp(x f)dx = f$.	(1) <i>calibration, reliability, or type 1 conditional bias</i> : $\mathbb{E}_f[f - \mathbb{E}(x f)]^2$; (2) <i>resolution</i> : $\mathbb{E}_f[\mathbb{E}(x f) - \mathbb{E}(x)]^2$.
$p(f)$	(related to <i>refinement, or sharpness</i>)	Refinement or sharpness is an aspect that usually applies only to probabilistic forecasts (Murphy, 1997). In deterministic forecast verification, if a forecaster produces the same forecast all the time, it is said to be completely unrefined. However, the complete refinement is difficult to define for deterministic forecasting (Murphy et al., 1989), but $p(f)$ has to be equal to $p(x)$ for perfect forecasts.	(1) <i>Kolmogorov–Smirnov test statistic</i> : $\max F(f) - F(x) $; (2) <i>earth mover’s distance or first Wasserstein distance</i> : the area between the two ECDFs. (The formal definition is technical and thus omitted.)
$p(f x)$	<i>likelihood</i>	If $p(f x)$ is zero for all values x but one, the forecast is perfectly discriminatory. If $p(f x)$ is the same for all values of x , the forecast is not at all discriminatory.	(1) <i>discrimination 1, or type 2 conditional bias</i> : $\mathbb{E}_x[x - \mathbb{E}(f x)]^2$; (2) <i>discrimination 2, or simply discrimination</i> : $\mathbb{E}_x[\mathbb{E}(f x) - \mathbb{E}(f)]^2$.
$p(x)$	<i>uncertainty, or base rate</i>	If $p(x)$ is a fairly peaked distribution, the scenario has relatively small uncertainty (and thus easier to forecast) as compared to a scenario where $p(x)$ is fairly uniform.	(1) <i>variance</i> : $\mathbb{V}(x)$; (2) <i>kurtosis</i> : $\frac{\mathbb{E}[(x - \mathbb{E}(x))^4]}{(\mathbb{E}[(x - \mathbb{E}(x))^2])^2}$.

is used in a solar forecasting context, it is intended to be applied to the clear-sky index, or other detrended variable of interest. However, the reference forecasts, i.e., r_i in Eq. (4), should be the original variable of interest, i.e., GHI, DNI, or PV power output. This is discussed further in Section 4.

3. Distribution-oriented approach for forecast verification

The distribution-oriented forecast verification framework is quite general. Before one starts to wonder what joint distribution¹⁵ has to do with forecast verification, most likely he or she has already used this framework. It is common to use a forecast–observation scatter plot to check forecast quality. One may draw some conclusions based on whether the point cloud is centered on the identity line, or how dispersed the scatter points are. In other cases, a forecaster may wish to check how the scatter points are distributed along the x-axis, or whether the spread of forecasts vary for different observation ranges. In fact, most forecast accuracy quantification—visually or through accuracy measures—are just summaries of the joint distributions, or equivalently, the marginal and conditional distributions. The relationship between joint, marginal, and conditional distributions of two random variables can be expressed using Bayes’ theorem. When these variables are the forecast and the observations, the same relationship is referred to as Murphy–Winkler factorization in meteorology (Murphy and Winkler, 1987).

Murphy–Winkler factorizations are:

$$p(f, x) = p(x|f)p(f), \quad (10)$$

$$p(f, x) = p(f|x)p(x), \quad (11)$$

where p denotes distribution, f and x represent forecasts and observations, respectively. Eq. (10) is called the *calibration–refinement factorization*, whereas Eq. (11) is called the *likelihood–base rate factorization*. The naming convention is quite intuitive. For example, the $p(x|f)$ term in Eq. (10) describes the spread of the observations, given a particular forecast. For a good correspondence, the forecast is said to be *calibrated* or *reliable*. Mathematically, the forecasts are perfectly calibrated if $\mathbb{E}(x|f) = \int xp(x|f)dx = f$. The reader is referred to Table 3 for an interpretation of other conditional and marginal distributions, and Murphy (1997) for a list of aspects of forecast quality.

Verifying the above conditional and marginal distributions is equivalent to verifying the joint distribution. For instance, given two sets of forecasts, f_1 and f_2 , by comparing $p(x|f_1)$ and $p(x|f_2)$, one can conclude whether one set of forecasts is more reliable than the other, see Moskaitis (2008); Murphy et al. (1989) for case studies. Whereas linking the forecast distributions to aspects of forecast quality provides forecasters with insights regarding their forecasts,

¹⁵Formally, we call a function $p(x, y)$ the joint distribution of random variables X and Y if $p(x, y) \geq 0, \forall(x, y); \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} p(x, y) dx dy = 1$; and for any set $\mathcal{A} \subset \mathbb{R} \times \mathbb{R}$, $\mathbb{P}[(X, Y) \in \mathcal{A}] = \iint_{\mathcal{A}} p(x, y) dx dy$, $\mathbb{P}[(X, Y) \in \mathcal{A}]$ denotes the probability of (X, Y) in set \mathcal{A} (Wasserman, 2013).

such results are easier to interpret if the different aspects of forecast quality can be quantified using measures. For instance, consider the bias–variance decomposition of MSE:

$$\begin{aligned}
\text{MSE} &= \iint (f - x)^2 p(f, x) df dx \\
&= \mathbb{E}[(f - x)^2] \\
&= \mathbb{V}(f - x) + [\mathbb{E}(f) - \mathbb{E}(x)]^2 \\
&= \underbrace{\mathbb{V}(f) + \mathbb{V}(x)}_{\text{marginal dist.}} - \underbrace{2\text{cov}(f, x)}_{\text{association}} + \underbrace{[\mathbb{E}(f) - \mathbb{E}(x)]^2}_{\text{unconditional bias}}, \tag{12}
\end{aligned}$$

where the overhead braces show the representation of each term. In this decomposition, $\mathbb{V}(f)$ and $\mathbb{V}(x)$ are variances of forecasts and observations, respectively. Their values can be used as a proxy for measuring the similarity between $p(f)$ and $p(x)$. If the forecasts were perfect, the two marginal distributions would be exactly the same, and so would the variances.¹⁶ Similarly, the $\text{cov}(f, x)$ term can be written as correlation, namely, $\sqrt{\mathbb{V}(f)\mathbb{V}(x)} \cdot \text{cor}(f, x)$, which denotes the *association* between forecasts and observations. Lastly, the $[\mathbb{E}(f) - \mathbb{E}(x)]^2$ term represents the squared unconditional bias, i.e., MBE². This example illustrates the complementarity between the measure-oriented approach (e.g., verification using MBE, correlation, or variance of the forecasts) and the distribution-oriented approach (analyzing the joint, conditional, and marginal distributions of forecasts and observations). Some of the widely used measures in solar forecasting, such as MBE, RMSE, or the Kolmogorov–Smirnov test integral, are related to their respective distributions in [Appendix C](#).

Besides the bias–variance decomposition, MSE can also be decomposed following the calibration–refinement and likelihood–base rate factorizations:

$$\text{MSE} = \mathbb{V}(x) + \underbrace{\mathbb{E}_f[f - \mathbb{E}(x|f)]^2}_{\text{type 1 conditional bias}} - \underbrace{\mathbb{E}_f[\mathbb{E}(x|f) - \mathbb{E}(x)]^2}_{\text{resolution}}, \tag{13}$$

$$\text{MSE} = \mathbb{V}(f) + \underbrace{\mathbb{E}_x[x - \mathbb{E}(f|x)]^2}_{\text{type 2 conditional bias}} - \underbrace{\mathbb{E}_x[\mathbb{E}(f|x) - \mathbb{E}(f)]^2}_{\text{discrimination}}. \tag{14}$$

A derivation of these decompositions is shown in [Moskaitis \(2008\)](#). As annotated above the equations, different terms in the decomposed forms explain different aspects of forecast quality.

The *type 1 conditional bias*, $\mathbb{E}_f[f - \mathbb{E}(x|f)]^2$, indicates the degree of correspondence between the mean observation given a particular forecast and the conditioning forecast, i.e., calibration or reliability. Recall that perfect calibration is when $\mathbb{E}(x|f) = f$, so the smaller this term is, the better. *Resolution* accounts for the difference between conditional and unconditional mean observation, which is reflected by $\mathbb{E}_f[\mathbb{E}(x|f) - \mathbb{E}(x)]^2$. If $\mathbb{E}(x|f) = \mathbb{E}(x)$, it means the data samples have no resolution. Since it is desirable to have the generated forecasts to be followed by different observations (so that the forecasts are meaningful), this term should be maximized. This is also reflected by the negative sign in front of that term. The *type 2 conditional bias*, $\mathbb{E}_x[x - \mathbb{E}(f|x)]^2$, indicates the degree of correspondence between the mean forecast given a particular observation and the observation. Naturally, this term should be as small as possible. Lastly, *discrimination* denotes the difference between the conditional and unconditional mean forecasts, i.e., $\mathbb{E}_x[\mathbb{E}(f|x) - \mathbb{E}(f)]^2$, which indicates how forecasts are differentiated for different observation values. This terms needs to be maximized.

The numerical evaluation of these decomposed factors can be difficult. When [Murphy and Winkler \(1987\)](#) proposed these decompositions, a binary x was used in their case study, which greatly simplifies the computation. In [Moskaitis \(2008\)](#), the evaluation was performed by discretizing the continuous random variable—tropical cyclone intensity—into bins. Recently, [Yang and Perez \(2019\)](#) used kernel conditional density estimation (KCDE) to estimate the conditional expectations, namely, $\mathbb{E}(x|f)$ and $\mathbb{E}(f|x)$, which removes the dependency on binning. The code for the KCDE-based approach is available in the supplementary material of that paper.

¹⁶However, having identical variances does not imply identical distributions; and having identical distributions does not imply the forecasts are perfect.




Figure 3: Joint and marginal distributions of 24-h-ahead hourly NAM forecasts and SURFRAD observations at (a) Desert Rock, Nevada (36.624°N, 116.019°W), and (b) Penn. State Univ., Pennsylvania (40.720°N, 77.931°W), from 2015 to 2016. The contour lines show the 2d kernel densities.

442 In contrast to numerical evaluation of Eqs. (13) and (14), visual inspection is more straightforward and enables a
 443 forecaster to appreciate the properties of the forecasts in great detail. In general, visualizing the error distribution is a
 444 powerful way of communicating the performance of a model. In line with the Murphy–Winkler factorizations, an x – f
 445 scatter plot displays the joint distribution between observations and forecasts, and allows for visualizing the marginal
 446 distributions as well as specific conditional distributions.

447 To exemplify the forecast verification procedure discussed in this section, a case study is presented. Fig. 3 shows
 448 the joint and marginal distributions of 24-h-ahead hourly forecasts of global horizontal irradiance (GHI) produced by the
 449 North American Mesoscale (NAM) forecast system against the observations collected by the Surface Radiation
 450 Budget Network (SURFRAD), at two locations with distinct climate over a period of two years. Whereas the joint
 451 distribution at the Desert Rock (DRA) station has approximately equal probabilities on both sides of the diagonal,
 452 the forecasts at the Penn. State Univ. (PSU) station over-predict GHI, i.e., the probability density is higher above
 453 the diagonal, where $f > x$. A closer examination of the 2d kernel density contours reveals that the NAM forecasts
 454 at DRA drift slightly below the identity line for high-irradiance conditions. For mid- and low-irradiance conditions
 455 at DRA, the forecasts are slightly above the identity line. This observation warrants an irradiance-condition-based
 456 post-processing treatment. Similar observations can be made for forecasts at PSU.

457 The histograms shown in Fig. 3 denote marginal distributions, $p(f)$ (to the right) and $p(x)$ (on the top). Since the
 458 shape of the histograms depends largely on bin width, different choices may affect the forecaster’s judgment differ-
 459 ently. In this regard, overlaying the empirical cumulative distribution functions (ECDFs) of f and x could be useful at
 460 times. Fig. 4 demonstrates such plots using the same data. Visually, the ECDFs of forecasts and observations at DRA
 461 align better than those at PSU. At PSU, f is stochastically greater than x (the ECDF of f lies below and hence to the
 462 right of that for x). Formally, the Kolmogorov–Smirnov test computes the statistic $D_n = \max |F_n(f) - F_n(x)|$, i.e.,
 463 the maximum absolute distance between the ECDFs of forecasts and observations. In the present case, Kolmogorov–
 464 Smirnov tests conducted at the two stations both reject the null hypothesis—two distributions are equal—at a signifi-
 465 cance level of 0.05.

466 As compared to joint and marginal distributions, the visualization of conditional distributions is more challenging.
 467 Whereas some authors plot the individual quantiles or use box plots to represent the distributions, ridgeline plots
 468 are employed here, see Fig. 5. In this plot, $p(x|f)$ and $p(f|x)$ at both stations are represented using overlapping lines,




Figure 4: Marginal distributions of forecasts and observations described in Fig. 3.

Table 4: Bias–variance decomposition (see Eq. 12) and Murphy–Winkler factorization (see, Eqs. 13 and 14) of 1–24-h-ahead NAM (f) against SURFRAD GHI (x), at DRA and PSU stations over 2015–2016. For interpretability, all metrics are written in squared form, so that the bases have the unit of W/m^2 , except for correlation ρ , which is dimensionless.

	MSE	$\mathbb{V}(x)$	$\mathbb{V}(f)$	$\rho(f, x)$	$[\mathbb{E}(f) - \mathbb{E}(x)]^2$	$\mathbb{E}_f[f - \hat{\mathbb{E}}(x f)]^2$	$\mathbb{E}_f[\hat{\mathbb{E}}(x f) - \mathbb{E}(x)]^2$	$\mathbb{E}_x[x_g - \hat{\mathbb{E}}(f x)]^2$	$\mathbb{E}_x[\hat{\mathbb{E}}(f x) - \mathbb{E}(f)]^2$
DRA	108.29 ²	297.70 ²	288.42 ²	0.94	22.67 ²	28.02 ²	278.68 ²	38.88 ²	270.09 ²
PSU	155.11 ²	269.83 ²	275.26 ²	0.85	41.73 ²	64.41 ²	229.86 ²	61.24 ²	235.40 ²

469 which create the impression of a mountain range. Fig. 5 (a) and (c) reveal that $p(x|f)$ is mostly centered on the forecast
470 value, i.e., $\mathbb{E}(x|f)$ is close to f , indicating small type 1 conditional bias in NAM forecasts. On the other hand, type 2
471 conditional bias is found to be significant for high values of x , see $p(f|x)$ for $x = 1050 \text{ W/m}^2$ in Fig. 5 (b) and (d).

472 The distribution-oriented verification technique is often complemented with *summary measures* of the different
473 aspects of forecast quality. Table 4 shows the quantification of these aspects using the bias–variance decomposition
474 and Murphy–Winkler factorization, as stated in Eqs. (12)–(14). Note that the decomposed terms listed in Table 4
475 do not add up exactly to MSE, due to the uncertainty introduced during KCDE. Such discrepancy is however small,
476 and thus does not affect this analysis. In terms of correlation, a higher $\rho = 0.94$ is observed at DRA as compared to
477 0.85 at PSU, indicating a better association between forecasts and observations at DRA. The square of unconditional
478 bias, $[\mathbb{E}(f) - \mathbb{E}(x)]^2$, is also significantly smaller at DRA, agreeing with the earlier observation made using the joint
479 distribution plots. Since a smaller type 1 conditional bias, $\mathbb{E}_f[f - \hat{\mathbb{E}}(x|f)]^2$, means higher calibration—the forecasts
480 at DRA are more reliable than those at PSU. Similarly, smaller type 2 conditional bias, higher resolution, and higher
481 discrimination observed at DRA all lead to the conclusion that the NAM forecasts at DRA have better forecast quality
482 than those at PSU. This conclusion confirms what could be expected because of the much cloudier climate of the latter
483 station.

484 Based on the case study above, it is evident that the distribution-oriented forecast verification is useful in assisting
485 forecasters to make informed decision based on forecast quality. In other cases, the same methodology can be applied
486 to compare forecasts made using different methods, thus providing more information than using MSE values alone.
487 This verification procedure leads to more meaningful conclusions than statements such as “the MSE at location A
488 is smaller than that at location B, and thus the forecasts at location A are better.” Nevertheless, should one wish to
489 examine the relative accuracy gain from the reference method, the quantification of aspects of forecast quality can
490 be carried out with the reference forecasts, namely, the optimal convex combination of climatology and persistence,
491 as exemplified in Table 5. Comparing the two tables, it is obvious that in the case of NAM, the NWP-based model
492 dominates the 24-h persistence in all aspects except for the unconditional and type 1 conditional bias. Nonetheless,
493 such bias in NWP forecasts can be trimmed with regression-based post-processing, and thus does not affect one’s
494 confidence in opting for the NAM model.

495 To promote the uptake of this distribution-oriented approach, data and code for reproducing all results appeared
496 in this section are provided in Appendix D.




Figure 5: Conditional distributions of 24-h-ahead hourly NAM forecasts and SURFRAD observations. $p(x|f)$ are shown in (a) and (c) for Desert Rock, Nevada (36.624°N, 116.019°W) and Penn. State Univ., Pennsylvania (40.720°N, 77.931°W), respectively. $p(f|x)$ are shown in (b) and (d) for the two stations, respectively.

497 4. Recommendations and practical concerns

498 The traditional measure-oriented approach is complemented by the distribution-oriented approach. The two ap-
 499 proaches reflect different aspects of forecast quality and help forecasters analyze their forecasts. Generally, forecasters
 500 are encouraged to use any meaningful measure to gauge their forecasts. Ultimately, however, a “one-number sum-
 501 mary” of forecasts is still highly desirable, especially when scientists bring their forecasts to non-technical personnel,
 502 e.g., sales persons, politicians, or the general public. Therefore, in this section, some recommendations and practical
 503 concerns regarding the use of skill scores are discussed.

Table 5: Same as Table 4, but tabulating the forecast quality of 24-h-ahead forecasts made using the optimal convex combination of climatology and persistence on clear-sky index.

	MSE	$\mathbb{V}(x)$	$\mathbb{V}(f)$	$\rho(f, x)$	$[\mathbb{E}(f) - \mathbb{E}(x)]^2$	$\mathbb{E}_f[f - \hat{\mathbb{E}}(x f)]^2$	$\mathbb{E}_f[\hat{\mathbb{E}}(x f) - \mathbb{E}(x)]^2$	$\mathbb{E}_x[x_g - \hat{\mathbb{E}}(f x)]^2$	$\mathbb{E}_x[\hat{\mathbb{E}}(f x) - \mathbb{E}(f)]^2$
DRA	115.94 ²	297.70 ²	259.82 ²	0.92	6.96 ²	22.33 ²	274.99 ²	60.12 ²	240.08 ²
PSU	168.06 ²	269.83 ²	185.12 ²	0.79	14.77 ²	37.19 ²	214.20 ²	126.84 ²	148.57 ²

4.1. MBE-, MAE-, or RMSE-based skill score?

The selection of RMSE to evaluate the skill score was made in Section 2.2.3. The arguments for that choice are presented in this section. The skill score belongs to the class of relative measures (Hyndman and Koehler, 2006),¹⁷ which means that a scale-dependent measure is needed for its computation. Since MBE, MAE, and RMSE are currently the most popular metrics (Yang et al., 2018), the discussion below considers all these options.

MBE is defined as $\mathbb{E}(f - x)$, or equivalently, $\mathbb{E}(f) - \mathbb{E}(x)$. An opposite definition, $\mathbb{E}(x - f)$, also exists in the statistics literature, which can be confusing. The latter definition originates from how a predictive model is constructed (see Makridakis et al., 2008).¹⁸ Defining the MBE to be “forecast minus observation” is more natural for solar forecasting, since an over-prediction (where forecasts are, on average, higher than observations) corresponds to a positive MBE, and an under-prediction corresponds to a negative MBE. MBE describes unconditional bias, and most statistical forecasting methods have $\text{MBE} \rightarrow 0$. State-of-the-art operational solar forecasts would have some form of bias correction implemented, e.g., model output statistics (MOS). Therefore, having small MBE is more of a baseline requirement, rather than a credit-worthy feature among state-of-the-art forecasts. Furthermore, the MBE of the reference forecasts from clear-sky persistence has an expectation of zero, and thus makes the skill score undefined. To that end, MBE is unsuitable for skill score computation.

The main difference between MAE, defined as $\mathbb{E}(|f - x|)$, and RMSE, defined as $\sqrt{\mathbb{E}[(f - x)^2]}$, is that the latter penalizes large errors while the former gives the same weight to all errors. Since large errors are particularly concerning for grid integration of solar power (e.g., a loss of load becomes more likely), RMSE is more suitable when a set of forecasts contain several large errors, which is usually the case for solar forecasts. From that standpoint, the percentage improvement in RMSE, in the form of s , might attract more interests than the MAE skill score.

The second reason for using the RMSE skill score is related to the distribution-oriented forecast verification. The Murphy–Winkler MSE factorization has been recommended for forecast analysis in the previous section. As a result, RMSE values become readily available after the various aspects of forecast quality are quantified. It must be highlighted that in the field of meteorology, and many other fields such as statistics, researchers generally do not distinguish “RMSE” and “MSE” in their writing (Jolliffe and Stephenson, 2012), simply because RMSE and MSE differ only by a square root. Nonetheless, they should not be mixed up during skill score computations, as evidenced by the example given in Section 2.2.3.

Lastly, the popularity of RMSE is higher than that of MAE, not only in solar forecasting, but in other forecasting domains as well. Gneiting (2011) found that the usage of RMSE in four related domains, namely, forecasting, statistics, econometrics, and meteorology, dominates as compared to MAE. Whereas the precise reasons are unknown, it is hypothesized that consistency might be one of the main reasons, since there are more models minimizing MSE than minimizing MAE. Additionally, squares are more amenable than absolute values in many mathematical operation (Chai and Draxler, 2014). Hence, for this and the above reasons, the RMSE skill score is recommended in deterministic solar forecasting.

4.2. Climatology, persistence, or their convex combination?

There are different definitions of what constitutes a climatology method. Murphy (1988) considered single-valued internal climatology, multiple-valued internal climatology, single-valued external climatology, and multiple-valued external climatology—each resulting in a different skill score expression. Whereas “single-valued” refers to methods

¹⁷One should distinguish relative measure from measure of relative error. The former performs division after the primary measure is computed, i.e., $\mathbb{E}[S(f, x)]/\mathbb{E}[S(r, x)]$, whereas the latter performs averaging on relative errors, i.e., $\mathbb{E}[S(f, x)/S(r, x)]$, where S is a scoring function.

¹⁸For a regression model, $y = g(x) + e$, where y is the response, x is the regressor, and bias e is in fact $\mathbb{E}[y - g(x)]$, or observation minus prediction.

that issue a single forecast for all occasions, “multiple-valued” methods issue different forecasts based on some conditional variables (e.g., one value for each season). The definitions of “internal” and “external” are based on whether the reference forecasts are derived directly from the n samples being verified, or derived from historical samples that are not included in the n samples. For instance, given an experimental period of one year, the sample mean, $\mathbb{E}(x)$, of all daytime observations in that year, is called the internal single-valued climatology forecast. Under this definition, the MSE of single-valued internal climatology is simply the sample variance, i.e.,

$$\text{MSE}_c = \mathbb{E} [\mathbb{E}(x) - x]^2 = \mathbb{V}(x). \quad (15)$$

539 It might be however argued that the multiple-valued climatology would be more appropriate when the forecasts
 540 of interest are to be evaluated over a time period longer than a month or a season (Murphy, 1988). For instance,
 541 such a climatology forecast is typically constructed for each month or each season. In the most extreme cases, the
 542 forecasts issued by multiple-valued climatology could be constructed in a rolling manner, with the averaging window
 543 equal to the forecast horizon. For example, in a day-ahead scenario of hourly GHI forecasts, one would use the
 544 averaged observed clear-sky index from the last available day as the “climatological” forecasts for each hour of the
 545 next day. This day-ahead reference method has been used by Beyer et al. (2009); Perez et al. (2013) among others,
 546 and has been experimentally validated by Yang (2019c). However, due to the many choices in constructing multiple-
 547 valued climatology forecasts, recommending the multiple-valued climatology defeats the purpose of standardizing
 548 the reference method used in solar forecasting. For this reason, the multiple-valued climatology will not be discussed
 549 further in this article.

On the other hand, Murphy (1992) defined the persistence forecast as a forecast based solely on the value of the variable of interest at an initial time. Given a forecast horizon, h , the MSE of h -step-ahead persistence can be written as:

$$\text{MSE}_p = \mathbb{E} (x_{i-h} - x_i)^2 = 2(1 - \gamma_h)\mathbb{V}(x), \quad (16)$$

550 where γ_h is the lag- h autocorrelation function. By comparing MSE_p to MSE_c , it is obvious that the relative perfor-
 551 mance of single-valued internal climatology and h -step-ahead persistence depends on γ_h . When $\gamma_h < 0.5$, $\text{MSE}_c <$
 552 MSE_p . When $\gamma_h > 0.5$, $\text{MSE}_c > \text{MSE}_p$. The two MSEs are equal if and only if $\gamma_h = 0.5$.

To ensure that the value of γ_h , which depends on the association among observations at lag h , does not affect a forecaster’s choice of reference method, one can consider a convex combination of the two. The combined forecast is $\alpha x_{i-h} + (1 - \alpha)\mathbb{E}(x)$. The optimal value of α can be obtained by differentiating the MSE of the climatology–persistence combination method, and equating it to zero, i.e., $d\text{MSE}_{cp}/d\alpha = 0$. It is straightforward to see that the optimal α is simply γ_h , i.e., $\alpha = \gamma_h$. The MSE of the optimal convex combination of single-valued internal climatology and h -step-ahead persistence is thus:

$$\begin{aligned} \text{MSE}_{cp} &= \mathbb{E} [\gamma_h x_{i-h} + (1 - \gamma_h)\mathbb{E}(x) - x_i]^2 \\ &= \mathbb{E} \{ \gamma_h [x_{i-h} - \mathbb{E}(x)] - [x_i - \mathbb{E}(x)] \}^2 \\ &= \gamma_h^2 \mathbb{V}(x) + \mathbb{V}(x) - 2\gamma_h \text{cov}(x_{i-h}, x_i) \\ &= (1 - \gamma_h^2) \mathbb{V}(x), \end{aligned} \quad (17)$$

where cov denotes the covariance. By comparing the expression of MSE_{cp} to MSE_c and MSE_p , it is apparent that $\text{MSE}_{cp} = (1 - \gamma_h^2)\text{MSE}_c$ and $\text{MSE}_{cp} = [(1 + \gamma_h)/2]\text{MSE}_p$. Since $-1 \leq \gamma_h \leq 1$, one obtains $\text{MSE}_{cp} \leq \text{MSE}_c$ and $\text{MSE}_{cp} \leq \text{MSE}_p$, with equality only when $\gamma_h = 0$ and $\gamma_h = 1$, respectively. Consequently, the RMSE of the climatology–persistence combination is always smaller or equal to that of either reference method, that is,

$$\text{RMSE}_{cp} \leq \min(\text{RMSE}_c, \text{RMSE}_p). \quad (18)$$

553 This condition is a powerful support for choosing the climatology–persistence combination as the standard of refer-
 554 ence method in deterministic solar forecasting. This choice has been discussed by Yang (2019b,c), with extensive
 555 empirical validation, for intra-hour, intra-day, and day-ahead scenarios.

4.3. Choice of the clear-sky model

Although it should be obvious at this point, it is re-emphasized that the naïve reference methods in solar forecasting always operate on clear-sky index. A clear-sky radiation model is thus needed to convert all-sky irradiance to clear-sky index and back. When developing forecasts using an NWP model, chances are the model can output the ideal clear-sky irradiance at regular time steps. When these direct outputs are not available, one of the many clear-sky radiation models that are available must be selected and run separately. Several extensive reviews on those models that are typically used in solar applications have been published recently (e.g., Sun et al., 2019; Antonanzas-Torres et al., 2019; Ruiz-Arias and Gueymard, 2018b). The best such models are simpler, but also faster, than those used in NWP models, while providing surface irradiance estimates of similar accuracy.

One particular issue is that most of the high-performance clear-sky models require several inputs, such as total column ozone amount, precipitable water, aerosol optical depth, or aerosol single-scattering albedo. Most generally, these quantities are not observed locally and must be sourced from remote-sensing or reanalysis databases. The process of obtaining such data at the proper temporal resolution from the highest-quality sources is not easy and may appear intimidating to non-experts. On the other hand, very simple radiation models that require only a few undemanding input variables usually have too limited performance for serious work (Gueymard, 2012). Therefore, the McClear model (Gschwind et al., 2019; Lefèvre et al., 2013) might be the best choice for solar forecasters, at least if they do not need to delve into the distant past (since the model delivers data only after 2004-01-01). Being a physical model based on radiative transfer, the performance of McClear is among the best (Sun et al., 2019; Ruiz-Arias and Gueymard, 2018a). McClear is conveniently available as a web service¹⁹ that delivers its predictions at five different time scales (1 min, 15 min, hourly, daily, and monthly) for all locations in the world, and for 2004-01-01 up to two days ago.²⁰ The R package “camsRad” is also freely available, and offers access to McClear through an API. The only limitation of McClear comes from its spatio-temporal resolution constraints, imposed by the reanalysis data it uses as its main inputs: 0.5° in both latitude and longitude on a 3-hourly basis. Hence, over complex terrain (e.g., mountains) or during fast-changing atmospheric conditions (e.g., dust storms), the McClear predictions might significantly differ from the actual situation at any specific site within its defining $0.5^\circ \times 0.5^\circ$ pixel. In such cases, the forecaster would have the burden of appropriately correcting the clear-sky predictions.

4.4. Some seemingly trivial implementation issues

Forecast verification must be based on appropriate out-of-sample tests (see Tashman, 2000, for a review on various test designs). There are several seemingly trivial implementation issues such as data trimming, normalization, counting nighttime hours, and data aggregation, that can strongly affect the verification results.

Data trimming refers to quality control (QC) applied to experimental data prior to forecasting and verification. Owing to factors such as measurement uncertainty or irradiance modeling error, experimental data often contain spurious values. There is no universally-accepted QC procedure (Gueymard and Ruiz-Arias, 2016), but recommended QC methods for surface radiation measurements (Long and Shi, 2008; García et al., 2014) and PV power output (Killinger et al., 2017) do exist. Recently, owing to the advances in remote-sensing technology, satellite-derived irradiance products have been shown to help perform QC on irradiance measurements (Urraca et al., 2017; Perez et al., 2017).

As statistical and machine-learning software packages become more powerful, details of implementation are now more opaque to forecasters. Therefore, forecasters should check the output of each step during forecasting, responsibly, to prevent spurious data from entering the final verification stage. To ensure forecast verification is performed with reasonably trimmed data, visual inspection as outlined in the previous section is necessary. It is however noted that data trimming based on forecast error is *not* recommended. One should not remove a forecast–observation pair just because it produces a large error; instead, the cause behind it should be investigated.

In parallel, there is also no well-accepted answer to the question “whether the nighttime hours should be included during validation.” Inclusion of nighttime hours would make the overall error smaller, since the forecasts— 0 W/m^2 —are perfect during those hours. Hence, nighttime data should always be excluded from verification. This could

¹⁹<http://www.soda-pro.com/web-services/radiation/cams-mcclear>

²⁰This two-day lag makes it unsuitable for operational forecast verification. Nonetheless, during forecast verification, one is only interested in analyzing the long-term behavior of different models, so that this lag may not be of concern.

602 be ensured by using a zenith angle filter of $< 85^\circ$ —a provision that also removes low-sun situations during which
603 measurements are less accurate and irradiance is too low to be of significance in solar power applications. Nonetheless,
604 special care is needed when using the filtered data. For instance, the sample average used in single-valued internal
605 climatology should be based on the filtered data, whereas the lag- h autocorrelation should be computed based on the
606 unfiltered data, to preserve the temporal spacing of the original clear-sky index time series.

607 In state-of-the-art solar forecasting, it is common to have more than one data source involved (e.g., ground mea-
608 surements, satellite-derived irradiance, reanalysis data, NWP output, or PV output). Even in the case of verifying
609 NWP forecasts only, one often desires to compare them to ground-based measurements. The issue of data aggregation
610 naturally comes into play because the ground-based measurements are usually at a higher temporal resolution (e.g., 1
611 min) than the NWP output (e.g., hourly). A related issue is caused by occurrences of data gaps in the data stream. This
612 is a relatively frequent issue, caused by either instrument malfunction or detection of incorrect or suspicious data by
613 the QC procedure. Data gaps create a difficulty when attempting to obtain hourly or daily averages. Some gap-filling
614 or statistical methods do exist, however, to limit the resulting uncertainty in the temporal means.

615 There are three schemes of averaging to aggregate a high-temporal-resolution time series into a low-temporal-
616 resolution one, namely, floor, ceiling, and round. A floor aggregation means that the data within a time interval are
617 aggregated to the earliest timestamp in that interval, e.g., 1-min data points between 1:00 to 2:00 are aggregated and
618 stamped as 1:00. Similarly, the ceiling-aggregation scheme stamps the aggregated data with the last timestamp of an
619 interval, and the round-aggregation scheme collapses the data to the center timestamp of an interval. Experimental
620 irradiance data may follow any one of these conventions (Polo et al., 2019), which complicates matters. It is obvious
621 that inappropriate data aggregation creates temporal misalignment between different datasets. This in turn may arti-
622 ficially amplify forecast errors (see Fig. 1 in Yang, 2018). To select the correct data-aggregation scheme, it is necessary
623 to understand how each dataset is produced. In other words, one must always read the data documentation. For
624 instance, most radiometric networks use the ceiling convention, with the notable exception of the Baseline Surface
625 Radiation Network (BSRN), which uses the floor convention. Satellite-derived irradiance and some NWP outputs
626 have a “snapshot” nature, and the round-aggregation scheme is appropriate. In the case of reanalysis, the data often
627 represent the condition over the past hour, and the ceiling-aggregation schemes is appropriate.

628 4.5. Irradiance to PV power output conversion issues

629 As discussed at the beginning, solar forecasting refers to both solar irradiance forecasting and solar power output
630 forecasting. Given the fact that most modern solar forecasting methods leverage heavily on exogenous input, such
631 as camera, satellite, or NWP data, time series methods are less frequently published (Yang, 2019a). In this regard,
632 most researchers would take a two-step procedure to forecast solar power output: (1) forecast solar irradiance, and (2)
633 convert the irradiance forecasts to power forecasts. For a fair comparison of forecasts, specific plant attributes need to
634 be considered.

635 In the case of PV power plants, the panels are usually installed on a surface tilted at an angle close to the site’s
636 latitude, or with single- or double-axis trackers, to maximize the annual electricity production. Consequently, the
637 variability of solar irradiance on the plane of array (POA) is often higher than that on a horizontal surface, due to
638 the larger swing of irradiance. On the other hand, the physical size of the PV power plant should also be considered,
639 because a PV plant behaves as a low-pass filter whose cut-off frequency depends on its size. In addition to this, due
640 to the well-known geographical smoothing effect (Lave et al., 2012), the total output of a utility-scale PV power plant
641 or a cluster of PV plants has less variability. Generally, areal forecasts are more accurate than the forecasts made for
642 a point location (Yang et al., 2017). The reader is referred to Lohmann (2018) for a short review on spatio-temporal
643 variability of solar irradiance, and Perpiñán et al. (2013); Marcos et al. (2012) for analyses on power fluctuations in a
644 large PV plant and a cluster of PV plants, respectively.

645 To convert the irradiance to PV power, three classes of models are often used, namely, separation models (Gueymard
646 and Ruiz-Arias, 2016), transposition models (Yang, 2016), and irradiance-to-power conversion models (see Fig. 1 in
647 Li et al., 2017). Separation models seek to estimate the DNI and DHI components from GHI; they are only needed if
648 both DNI and DHI are unavailable. In the recent work by Yang and Boland (2019), several state-of-the-art separation
649 models were compared, and the Yang2 model was shown to be one of the most accurate separation models. This
650 result was later confirmed by Blaga (2019). On the other hand, transposition models convert horizontal irradiance
651 components to the POA irradiances, and the Perez1990 model (Perez et al., 1990) is commonly regarded as one of the
652 best-performing transposition models (Yang, 2016). Lastly, based on the POA irradiance and other parameters such

653 as ambient or module temperature, the power output from a PV plant with known attributes can be predicted fairly
654 accurately.

655 Although these modeling steps usually have a much higher accuracy than forecasting, their uncertainties have to
656 be considered during forecast verification. Generally, due to cancellation of errors, the overall error of PV power
657 forecasting is almost always smaller than the algebraic sum of errors in each step (forecasting, separation, transposi-
658 tion, and irradiance-to-power conversion). It is therefore of interest to study such error propagation, see [Urraca et al.
659 \(2018\)](#); [Almeida et al. \(2017\)](#) for some preliminary findings. With the verification framework discussed earlier, the
660 Murphy–Winkler factorizations could be applied to each of the steps, and thus provide new insights to solar engineers.

661 Lastly, it is worth noting that the baseline method for irradiance-to-power conversion is widely available in soft-
662 ware packages, such as the System Advisor Model, PVSyst, or pvlib. These software packages take irradiance time
663 series (e.g., typical meteorological year data) as input, and output PV power time series using the system specifications
664 provided by users. For forecast verification purposes, one can feed the forecast irradiance time series, and obtain the
665 corresponding forecast power time series. That said, emerging issues, such as lack of spectral and angular details in
666 irradiance input ([Lindsay et al., 2020](#)) or the effect of cloud-enhancement phenomenon on PV power plants ([Järvelä
667 et al., 2020](#)), are not yet being considered in these software packages, leaving room for improvements.

668 5. Conclusion

669 The increasing amount of solar forecasting research calls for harmonization of forecast verification measures
670 and methods among researchers. This paper has discussed a wide spectrum of issues relevant to verification of
671 deterministic solar forecasts. The final recommendations are listed as follows.

- 672 • The distribution-oriented approach to forecast verification can be used for forecast analysis. Since the joint
673 distribution contains all time-independent information relevant to verification, it is more general than the tradi-
674 tional measure-oriented approach. It is recommended to use the distribution-oriented approach to visualize and
675 quantify forecast quality.
- 676 • Bias–variance factorization and Murphy–Winkler factorizations link various qualitative aspects of the skillful-
677 ness of the forecasts (such as uncertainty, reliability, resolution, association, or discrimination) to quantitative
678 measures. These decomposed measures provide forecasters with realistic and insightful assessments of the
679 forecast verification problem, which essentially has a multifaceted nature.
- 680 • Small MBE is a prerequisite of all solar forecasts and therefore not a critical metric to judge forecast quality.
- 681 • When the normalized errors are reported, it is necessary to also tabulate the normalization values, such as the
682 mean irradiance values.
- 683 • Generally, forecasters are encouraged to use any meaningful measure to gauge their forecasts. However, if a
684 chosen measure is inconsistent with the given forecast directive, it might be inappropriate in theory.
- 685 • The RMSE skill score based on the optimal convex combination of single-valued internal climatology and
686 h -step-ahead clear-sky persistence is strongly recommended in all solar forecasting studies. The skill score
687 denotes the relative improvement of a model of interest from the reference method, and it can be used to
688 compare forecasts produced in different works. Nonetheless, the skill score only provides an overall idea, so
689 that the distribution-oriented approach is still required.
- 690 • Implementation issues are important for the final interpretation of forecast accuracy. Nighttime data should be
691 excluded from forecast verification. Special care is needed during data trimming and aggregation, as a result of
692 the necessary quality control and data gap-filling processes.
- 693 • As forecasting workflows are getting more and more complex, it is advised to perform sanity checks throughout
694 the course of producing forecasts. To ensure the worldwide uptake of any proposed forecasting model, source
695 code and data should be made available whenever possible. Without reproducibility, it would be cumbersome—
696 if not impossible—to verify the reported forecast performance.

697 **Acknowledgment**

698 The lead author would like to thank Cong Feng (University of Texas at Dallas), Bri-Mathias Hodge (National
699 Renewable Energy Laboratory), and Pierre Pinson (Technical University of Denmark) for their discussions during
700 manuscript preparation.

701 **Appendix A. On seasonal adjustment of reference forecasts**

702 In the simplest case, persistence, by definition, uses the most recent observation available as forecasts for all
703 horizons. Such forecast is also sometimes referred to as “random walk” (i.e., $r_t = r_{t-1} + e_t$), or “no change” forecast.
704 In the case of solar irradiance forecasting, raw persistence forecasts should be the most recent observed irradiance.
705 Notwithstanding, given the bell-shaped diurnal transient of irradiance caused by the Sun’s apparent movement, it
706 is important to take such seasonality into consideration. Makridakis et al. (2008) noted that seasonally adjusted
707 persistence can frequently do much better than the raw persistence. So the question is “how is it best to adjust for
708 seasonality?”

709 Besides using the clear-sky irradiance, one can use the extraterrestrial irradiance (the irradiance at the top of the
710 atmosphere) for adjustment. The ratio between surface and extraterrestrial GHI is known as clearness index, k . In
711 other words, the persistence is performed on clearness index, namely, $r_{t+h} = x_t \cdot k_{t+h} / k_t$. Both clear-sky persistence and
712 clearness persistence adopt a multiplicative seasonality modeling approach. A particular problem with multiplicative
713 seasonality is that during sunrise and sunset (small solar elevation angle), both the clear-sky index and clearness index
714 can become quite large, owing to the measurement uncertainty and the inaccuracy in the clear-sky models, and thus
715 the forecast errors can be large at those times. To exclude those undesirable forecasts that may severely distort the
716 error metrics, solar forecasters usually apply a zenith angle filter, e.g., zenith angle $< 85^\circ$, before error computation.
717 Alternatively, one can opt for an additive seasonality modeling, e.g., $r_{t+h} = x_t - c_t + c_{t+h}$. However, the remainder
718 series (i.e., $x_t - c_t$) in this case is still heteroscedastic.

719 Alternatively, one can also use a “cloudiness index” where the reference forecast is referred to as “smart persis-
720 tence” by Inman et al. (2015). This reference model includes the effects of air mass, aerosols, turbidity, i.e., every
721 major atmospheric effect except that of clouds. Because the timescale for turbidity variations is much larger than the
722 timescale for cloud-cover variations, the cloudiness index provides an excellent reference for short-time forecasts.
723 Any skill calculated over the cloudiness-index persistence measures the ability to capture cloud cover changes over
724 short periods of time, and thus the qualifier “smart”. Note that skills reported over smart persistence are necessarily
725 lower than skills reported over clear-sky persistence. This is because it is virtually impossible for a forecast to im-
726 prove over smart persistence under cloudless skies, since the smart persistence reference model includes all effects of
727 diurnal variability (solar zenith angle) and columnar optical depths of water vapor and aerosols. These types of smart
728 persistence reference forecasts are typically updated sub-hourly (Inman et al., 2015; Reno and Hansen, 2016).

729 The literature on the treatment of seasonal and multi-seasonal (e.g., diurnal and yearly cycles in solar irradi-
730 ance) time series is rich. Chapter 3 of Makridakis et al. (2008) provides details on various statistical techniques for
731 time series decomposition. In the solar forecasting literature, various techniques, such as Fourier series, exponential
732 smoothing, STL decomposition, additive clear-sky decomposition, or smoothed clear-sky decomposition, have also
733 been extensively explored (e.g., Dong et al., 2013; Yang et al., 2015; Voyant and Notton, 2018). However, as compared
734 to the clear-sky persistence, these methods usually require more steps, which may be a reason for their limited uptake.
735 Since the main goal of seasonally-adjusted persistence is to construct a better reference model than raw persistence,
736 clear-sky persistence offers a good trade-off between implementation difficulty and baseline accuracy.

737 **Appendix B. On equivalence of s^* and s**

Denoting the clear-sky index forecasts of interest, reference forecasts, and observations using ϕ , ρ , and κ , respec-
tively, one can write: $\phi_t \equiv f_t/c_t$, $\rho_t \equiv r_t/c_t$, and $\kappa_t \equiv x_t/c_t$. For 1-step-ahead clear-sky persistence, $r_t = x_{t-1} \cdot c_t/c_{t-1}$,

hence, $\rho_t = x_{t-1}/c_{t-1}$. Eqs. (8) and (9) can then be written as:

$$U = \sqrt{\frac{1}{N} \sum_{t=1}^N (\phi_t - \kappa_t)^2} = \text{RMSE}(\phi, \kappa), \quad (\text{B.1})$$

$$V = \sqrt{\frac{1}{N} \sum_{t=1}^N (\rho_t - \kappa_t)^2} = \text{RMSE}(\rho, \kappa). \quad (\text{B.2})$$

738 Therefore, $s^* = 1 - \text{RMSE}(\phi, \kappa)/\text{RMSE}(\rho, \kappa)$ is the RMSE skill score of clear-sky index forecasts, with 1-step-ahead
739 persistence reference.

Suppose a perfect clear-sky model exists, which explains all influences of atmospheric laws but those from clouds, then one can say that the clear-sky index is independent of the clear-sky expectation, i.e., $\kappa \perp\!\!\!\perp c$. Hence, any forecast error, $e = \phi - \kappa = g(\kappa) - \kappa$, where g is a function to produce κ forecast, is also independent of c , i.e., $e \perp\!\!\!\perp c$. From the properties of independence, it immediately follows that:

$$\mathbb{E}(e^2 c^2) = \mathbb{E}(c^2) \mathbb{E}(e^2). \quad (\text{B.3})$$

That is:

$$\text{MSE}(f, x) = \mathbb{E}(c^2) \text{MSE}(\phi, \kappa). \quad (\text{B.4})$$

Stated differently, if $\kappa \perp\!\!\!\perp c$, the MSE of clear-sky index forecasts and that of irradiance forecasts are scaled by a factor of $\mathbb{E}(c^2)$. Equivalently,

$$\text{RMSE}(f, x) = \sqrt{\mathbb{E}(c^2)} \cdot \text{RMSE}(\phi, \kappa). \quad (\text{B.5})$$

Since g is any function, the result also applies to reference κ forecast. Then,

$$s^* = 1 - \frac{\text{RMSE}(\phi, \kappa)}{\text{RMSE}(\rho, \kappa)} = 1 - \frac{\sqrt{\mathbb{E}(c^2)} \cdot \text{RMSE}(\phi, \kappa)}{\sqrt{\mathbb{E}(c^2)} \cdot \text{RMSE}(\rho, \kappa)} = 1 - \frac{\text{RMSE}(f, x)}{\text{RMSE}(r, x)} = s. \quad (\text{B.6})$$

740 That said, it should be clearly noted that the assumption $\kappa \perp\!\!\!\perp c$ is *not* valid in almost all cases. This is because
741 even the best clear-sky models could not completely explain the atmospheric effects on radiation. Hence, when there
742 is some degree of dependence between κ and c , one can only observe $s^* \approx s$, as empirically shown by [Marquez and](#)
743 [Coimbra \(2011, 2013\)](#); [Coimbra et al. \(2013\)](#). This issue deserves future attention.

744 Appendix C. Links between measure-oriented and distribution-oriented forecast verification approaches

The *rule of the lazy statistician* states ([Wasserman, 2013](#)): Let $y = g(x)$, then

$$\mathbb{E}(y) = \mathbb{E}[g(x)] = \int g(x) dF(x) = \int g(x) p(x) dx. \quad (\text{C.1})$$

The two-variable case is handled in a similar way: Let $z = g(x, y)$, then

$$\mathbb{E}(z) = \mathbb{E}[g(x, y)] = \int \int g(x, y) dF(x, y) = \iint g(x, y) p(x, y) dx dy. \quad (\text{C.2})$$

This rule links the joint distribution to a large collection of error metrics. For example, MBE, MAE, and RMSE can be written as:

$$\text{MBE} = \mathbb{E}[(f - x)] = \iint (f - x)p(f, x)dfdx, \quad (\text{C.3})$$

$$\text{MAE} = \mathbb{E}(|f - x|) = \iint |f - x|p(f, x)dfdx, \quad (\text{C.4})$$

$$\text{RMSE} = \sqrt{\mathbb{E}[(f - x)^2]} = \left[\iint (f - x)^2 p(f, x)dfdx \right]^{\frac{1}{2}}. \quad (\text{C.5})$$

Similarly, it is possible to express nMBE, nMAE, nRMSE, maximum absolute error, mean average percentage error, etc., in this form. Hence, it is clear that all these metrics are just different ways to summarize the joint distribution.

In the report by Beyer et al. (2009), four metrics based on the Kolmogorov–Smirnov test were proposed, namely, the Kolmogorov–Smirnov test integral (KSI), the OVER index, KSE (linear combination of KSI and OVER), and RIO (sum of KSD and RMSE, divided by 2). For instance, KSI calculates approximately the area between the ECDFs of f and x , whereas OVER calculates the area of those instances between the two ECDFs that exceed the critical value at 99% level of confidence. While the definitions of these metrics might deviate from the usual statistical measures (such as the Wasserstein distance), they actually constitute an early attempt by solar engineers to summarize the differences between marginal distributions of f and x .

Lastly, Section 3 demonstrates several ways to summarize the conditional distributions. In these cases, the summaries would also require marginal distributions. More specifically, the type 1 conditional bias and resolution are summaries of $p(x|f)$ and $p(f)$, whereas the type 2 conditional bias and discrimination are summaries of $p(f|x)$ and $p(x)$ (Murphy, 1997).

The remaining open questions, such as “are there better ways to summarize these distributions,” “which summaries allow cross-work forecast comparison,” “how to analyze the summaries graphically,” etc., jointly motivate future research on verification of deterministic solar forecasts.

Appendix D. Supplementary materials

Supplementary data associated with this article can be found, in the online version, at *doi place holder*

References

- Almeida, M.P., Muñoz, M., de la Parra, I., Perpiñán, O., 2017. Comparative study of PV power forecast using parametric and nonparametric PV models. *Solar Energy* 155, 854 – 866. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17306175>, doi:<https://doi.org/10.1016/j.solener.2017.07.032>.
- Antonanzas, J., Osorio, N., Escobar, R., Urraca, R., de Pison, F.M., Antonanzas-Torres, F., 2016. Review of photovoltaic power forecasting. *Solar Energy* 136, 78 – 111. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X1630250X>, doi:<https://doi.org/10.1016/j.solener.2016.06.069>.
- Antonanzas, J., Pozo-Vázquez, D., Fernandez-Jimenez, L., de Pison, F.M., 2017. The value of day-ahead forecasting for photovoltaics in the Spanish electricity market. *Solar Energy* 158, 140 – 146. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17308307>, doi:<https://doi.org/10.1016/j.solener.2017.09.043>.
- Antonanzas-Torres, F., Urraca, R., Polo, J., Perpiñán-Lamigueiro, O., Escobar, R., 2019. Clear sky solar irradiance models: A review of seventy models. *Renewable and Sustainable Energy Reviews* 107, 374 – 387. URL: <http://www.sciencedirect.com/science/article/pii/S1364032119301261>, doi:<https://doi.org/10.1016/j.rser.2019.02.032>.
- Armstrong, J.S., 2001. Evaluating forecasting methods, in: *Principles of forecasting*. Springer, pp. 443 – 472.
- Beyer, H.G., Polo Martinez, J., Suri, M., Torres, J.L., Lorenz, E., Müller, S.C., Hoyer-Klick, C., Ineichen, P., 2009. Benchmarking of Radiation Products. Technical Report 038665. Mesor Report D.1.1.3.
- Blaga, R., 2019. The impact of temporal smoothing on the accuracy of separation models. *Solar Energy* 191, 371 – 381. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X19308709>, doi:<https://doi.org/10.1016/j.solener.2019.08.078>.
- Blaga, R., Săbăduș, A., Ștefu, N., Dughir, C., Paulescu, M., Bădescu, V., 2019. A current perspective on the accuracy of incoming solar energy forecasting. *Progress in Energy and Combustion Science* 70, 119 – 144. URL: <http://www.sciencedirect.com/science/article/pii/S0360128518300303>, doi:<https://doi.org/10.1016/j.pecs.2018.10.003>.
- Chai, T., Draxler, R.R., 2014. Root mean square error (RMSE) or mean absolute error (MAE)? – Arguments against avoiding RMSE in the literature. *Geoscientific Model Development* 7, 1247 – 1250. URL: <https://www.geosci-model-dev.net/7/1247/2014/>, doi:<https://doi.org/10.5194/gmd-7-1247-2014>.

- 787 Coimbra, C.F.M., Kleissl, J., Marquez, R., 2013. Chapter 8 - Overview of solar-forecasting methods and a metric for accu-
788 racy evaluation, in: Kleissl, J. (Ed.), *Solar Energy Forecasting and Resource Assessment*. Academic Press, Boston, pp. 171 –
789 194. URL: <http://www.sciencedirect.com/science/article/pii/B9780123971777000085>, doi:[https://doi.org/10.1016/](https://doi.org/10.1016/B978-0-12-397177-7.00008-5)
790 [B978-0-12-397177-7.00008-5](https://doi.org/10.1016/B978-0-12-397177-7.00008-5).
- 791 Dong, Z., Yang, D., Reindl, T., Walsh, W.M., 2013. Short-term solar irradiance forecasting using exponential smoothing state space model. *Energy*
792 55, 1104 – 1113. URL: <http://www.sciencedirect.com/science/article/pii/S0360544213003381>, doi:[https://doi.org/10.](https://doi.org/10.1016/j.energy.2013.04.027)
793 [1016/j.energy.2013.04.027](https://doi.org/10.1016/j.energy.2013.04.027).
- 794 Fildes, R., Nikolopoulos, K., Crone, S.F., Syntetos, A.A., 2008. Forecasting and operational research: a review. *Journal of the Operational Research*
795 *Society* 59, 1150 – 1172. doi:<https://doi.org/10.1057/palgrave.jors.2602597>.
- 796 García, R.D., García, O.E., Cuevas, E., Cachorro, V.E., Romero-Campos, P.M., Ramos, R., de Frutos, A.M., 2014. Solar radiation measurements
797 compared to simulations at the BSRN Izaña station. mineral dust radiative forcing and efficiency study. *Journal of Geophysical Research:*
798 *Atmospheres* 119, 179–194. URL: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013JD020301>, doi:[https://](https://doi.org/10.1002/2013JD020301)
799 doi.org/10.1002/2013JD020301.
- 800 Gilleland, E., Ahijevych, D.A., Brown, B.G., Ebert, E.E., 2010. Verifying forecasts spatially. *Bulletin of the American Meteorological Society* 91,
801 1365 – 1376. doi:<https://doi.org/10.1175/2010BAMS2819.1>.
- 802 Gneiting, T., 2011. Making and evaluating point forecasts. *Journal of the American Statistical Association* 106, 746 – 762. URL: [http:](http://www.jstor.org/stable/41416407)
803 [//www.jstor.org/stable/41416407](http://www.jstor.org/stable/41416407), doi:<https://doi.org/10.2307/41416407>.
- 804 Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102,
805 359 – 378. doi:<https://doi.org/10.1198/016214506000001437>.
- 806 Gschwind, B., Wald, L., Blanc, P., Lefèvre, M., Schroedter-Homscheidt, M., Arola, A., 2019. Improving the McClear model estimating the
807 downwelling solar radiation at ground level in cloud-free conditions – McClear-v3. *Meteorologische Zeitschrift* 28, 147 – 163. doi:[https:](https://doi.org/10.1127/metz/2019/0946)
808 [//doi.org/10.1127/metz/2019/0946](https://doi.org/10.1127/metz/2019/0946).
- 809 Gueymard, C.A., 2012. Clear-sky irradiance predictions for solar resource mapping and large-scale applications: Improved validation methodology
810 and detailed performance analysis of 18 broadband radiative models. *Solar Energy* 86, 2145 – 2169. URL: [http://www.sciencedirect.](http://www.sciencedirect.com/science/article/pii/S0038092X11004221)
811 [com/science/article/pii/S0038092X11004221](http://www.sciencedirect.com/science/article/pii/S0038092X11004221), doi:<https://doi.org/10.1016/j.solener.2011.11.011>. progress in Solar En-
812 [ergy 3](https://doi.org/10.1016/j.solener.2011.11.011).
- 813 Gueymard, C.A., Ruiz-Arias, J.A., 2016. Extensive worldwide validation and climate sensitivity analysis of direct irradiance predictions from 1-
814 min global irradiance. *Solar Energy* 128, 1 – 30. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X15005435>,
815 doi:<https://doi.org/10.1016/j.solener.2015.10.010>. special issue: Progress in Solar Energy.
- 816 Hoff, T.E., Perez, R., Kleissl, J., Renne, D., Stein, J., 2013. Reporting of irradiance modeling relative prediction errors. *Progress in Photovoltaics:*
817 *Research and Applications* 21, 1514 – 1519. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.2225>, doi:[https:](https://doi.org/10.1002/pip.2225)
818 [//doi.org/10.1002/pip.2225](https://doi.org/10.1002/pip.2225).
- 819 Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., Hyndman, R.J., 2016. Probabilistic energy forecasting: Global Energy Forecasting
820 Competition 2014 and beyond. *International Journal of Forecasting* 32, 896 – 913. URL: [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0169207016000133)
821 [article/pii/S0169207016000133](http://www.sciencedirect.com/science/article/pii/S0169207016000133), doi:<https://doi.org/10.1016/j.ijforecast.2016.02.001>.
- 822 Huang, J., Thatcher, M., 2017. Assessing the value of simulated regional weather variability in solar forecasting using numerical weather prediction.
823 *Solar Energy* 144, 529 – 539. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17300774>, doi:[https://](https://doi.org/10.1016/j.solener.2017.01.058)
824 doi.org/10.1016/j.solener.2017.01.058.
- 825 Hyndman, R.J., Koehler, A.B., 2006. Another look at measures of forecast accuracy. *International Journal of Forecasting* 22, 679
826 – 688. URL: <http://www.sciencedirect.com/science/article/pii/S0169207006000239>, doi:[https://doi.org/10.1016/j.](https://doi.org/10.1016/j.ijforecast.2006.03.001)
827 [ijforecast.2006.03.001](https://doi.org/10.1016/j.ijforecast.2006.03.001).
- 828 Inman, R.H., Edson, J.G., Coimbra, C.F.M., 2015. Impact of local broadband turbidity estimation on forecasting of clear sky direct normal
829 irradiance. *Solar Energy* 117, 125 – 138. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X15002200>,
830 doi:<https://doi.org/10.1016/j.solener.2015.04.032>.
- 831 Inman, R.H., Pedro, H.T.C., Coimbra, C.F.M., 2013. Solar forecasting methods for renewable energy integration. *Progress in Energy and*
832 *Combustion Science* 39, 535 – 576. URL: <http://www.sciencedirect.com/science/article/pii/S0360128513000294>, doi:[https:](https://doi.org/10.1016/j.pecs.2013.06.002)
833 [//doi.org/10.1016/j.pecs.2013.06.002](https://doi.org/10.1016/j.pecs.2013.06.002).
- 834 Järvelä, M., Lappalainen, K., Valkealahti, S., 2020. Characteristics of the cloud enhancement phenomenon and PV power plants. *Solar Energy* 196,
835 137 – 145. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X19311909>, doi:[https://doi.org/10.1016/j.](https://doi.org/10.1016/j.solener.2019.11.090)
836 [solener.2019.11.090](https://doi.org/10.1016/j.solener.2019.11.090).
- 837 Jolliffe, I.T., 2008. The impenetrable hedge: a note on propriety, equitability and consistency. *Meteorological Applications* 15, 25 – 29. URL:
838 <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/met.60>, doi:<https://doi.org/10.1002/met.60>.
- 839 Jolliffe, I.T., Stephenson, D.B., 2012. *Forecast verification: a practitioner's guide in atmospheric science*. John Wiley & Sons.
- 840 Killinger, S., Engerer, N., Müller, B., 2017. QCPV: A quality control algorithm for distributed photovoltaic array power output. *Solar Energy* 143,
841 120 – 131.
- 842 Klingler, A.L., Teichtmann, L., 2017. Impacts of a forecast-based operation strategy for grid-connected PV storage systems on profitability and the
843 energy system. *Solar Energy* 158, 861 – 868. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17309325>,
844 doi:<https://doi.org/10.1016/j.solener.2017.10.052>.
- 845 Lave, M., Kleissl, J., Arias-Castro, E., 2012. High-frequency irradiance fluctuations and geographic smoothing. *Solar Energy* 86, 2190
846 – 2199. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X11002611>, doi:[https://doi.org/10.1016/j.](https://doi.org/10.1016/j.solener.2011.06.031)
847 [solener.2011.06.031](https://doi.org/10.1016/j.solener.2011.06.031). progress in Solar Energy 3.
- 848 Law, E.W., Kay, M., Taylor, R.A., 2016. Calculating the financial value of a concentrated solar thermal plant operated using direct normal irra-
849 diance forecasts. *Solar Energy* 125, 267 – 281. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X15007045>,
850 doi:<https://doi.org/10.1016/j.solener.2015.12.031>.
- 851 Lefèvre, M., Oumbe, A., Blanc, P., Espinar, B., Gschwind, B., Qu, Z., Wald, L., Schroedter-Homscheidt, M., Hoyer-Klick, C., Arola, A., Benedetti,

- 852 A., Kaiser, J.W., Morcrette, J.J., 2013. McClear: a new model estimating downwelling solar radiation at ground level in clear-sky conditions.
853 Atmospheric Measurement Techniques 6, 2403 – 2418. URL: <https://www.atmos-meas-tech.net/6/2403/2013/>, doi:<http://dx.doi.org/10.5194/amt-6-2403-2013>.
- 854
- 855 Li, Y., Chen, X., Zhao, B., Zhao, Z., Wang, R., 2017. Development of a PV performance model for power output simulation at minutely resolution.
856 Renewable Energy 111, 732 – 739. URL: <http://www.sciencedirect.com/science/article/pii/S0960148117303567>, doi:<https://doi.org/10.1016/j.renene.2017.04.049>.
- 857
- 858 Lindsay, N., Libois, Q., Badosa, J., Migan-Dubois, A., Bourdin, V., 2020. Errors in PV power modelling due to the lack of spectral and angular
859 details of solar irradiance inputs. Solar Energy 197, 266 – 278. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X19312563>, doi:<https://doi.org/10.1016/j.solener.2019.12.042>.
- 860
- 861 Lohmann, G.M., 2018. Irradiance variability quantification and small-scale averaging in space and time: A short review. Atmosphere 9. URL:
862 <https://www.mdpi.com/2073-4433/9/7/264>, doi:<https://doi.org/10.3390/atmos9070264>.
- 863
- 864 Long, C.N., Shi, Y., 2008. An automated quality assessment and control algorithm for surface radiation measurements. The Open Atmospheric
865 Science Journal 2, 23 – 37.
- 866
- 867 Lorenz, E., Kühnert, J., Heinemann, D., Nielsen, K.P., Remund, J., Müller, S.C., 2016. Comparison of global horizontal irradiance forecasts based
868 on numerical weather prediction models with different spatio-temporal resolutions. Progress in Photovoltaics: Research and Applications 24,
869 1626 – 1640. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.2799>, doi:[10.1002/pip.2799](https://doi.org/10.1002/pip.2799).
- 870
- 871 Madsen, H., Pinson, P., Kariniotakis, G., Nielsen, H.A., Nielsen, T.S., 2005. Standardizing the performance evaluation of short-term wind power
872 prediction models. Wind Engineering 29, 475 – 489. doi:<https://doi.org/10.1260/030952405776234599>.
- 873
- 874 Makridakis, S., Wheelwright, S.C., Hyndman, R.J., 2008. Forecasting methods and applications. John Wiley & sons.
- 875
- 876 Marcos, J., Marroyo, L., Lorenzo, E., García, M., 2012. Smoothing of PV power fluctuations by geographical dispersion. Progress in Photovoltaics:
877 Research and Applications 20, 226–237. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.1127>, doi:<https://doi.org/10.1002/pip.1127>.
- 878
- 879 Marquez, R., Coimbra, C.F.M., 2011. A novel metric for evaluation of solar forecasting models, in: ASME 2011 5th International Conference on
880 Energy Sustainability, ASME. pp. 1459 – 1467. doi:<https://doi.org/10.1115/ES2011-54519>.
- 881
- 882 Marquez, R., Coimbra, C.F.M., 2013. Proposed metric for evaluation of solar forecasting models. Journal of solar energy engineering 135, 011016.
883 doi:<https://doi.org/10.1115/1.4007496>.
- 884
- 885 Martinez-Anido, C.B., Botor, B., Florita, A.R., Draxl, C., Lu, S., Hamann, H.F., Hodge, B.M., 2016. The value of day-ahead solar power forecasting
886 improvement. Solar Energy 129, 192 – 203. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X16000736>,
887 doi:<https://doi.org/10.1016/j.solener.2016.01.049>.
- 888
- 889 van der Meer, D., Widén, J., Munkhammar, J., 2018. Review on probabilistic forecasting of photovoltaic power production and electricity con-
890 sumption. Renewable and Sustainable Energy Reviews 81, 1484 – 1512. URL: <http://www.sciencedirect.com/science/article/pii/S1364032117308523>, doi:<https://doi.org/10.1016/j.rser.2017.05.212>.
- 891
- 892 Moskaitis, J.R., 2008. A case study of deterministic forecast verification: Tropical cyclone intensity. Weather and Forecasting 23, 1195 – 1220.
893 doi:<https://doi.org/10.1175/2008WAF2222133.1>.
- 894
- 895 Murphy, A.H., 1988. Skill scores based on the mean square error and their relationships to the correlation coefficient. Monthly Weather Review
896 116, 2417 – 2424. doi:[https://doi.org/10.1175/1520-0493\(1988\)116<2417:SSBOTM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1988)116<2417:SSBOTM>2.0.CO;2).
- 897
- 898 Murphy, A.H., 1992. Climatology, persistence, and their linear combination as standards of reference in skill scores. Weather and Forecasting 7,
899 692 – 698. doi:[https://doi.org/10.1175/1520-0434\(1992\)007<0692:CPATLC>2.0.CO;2](https://doi.org/10.1175/1520-0434(1992)007<0692:CPATLC>2.0.CO;2).
- 900
- 901 Murphy, A.H., 1993. What is a good forecast? An essay on the nature of goodness in weather forecasting. Weather and Forecasting 8, 281 – 293.
902 doi:[https://doi.org/10.1175/1520-0434\(1993\)008<0281:WIAGFA>2.0.CO;2](https://doi.org/10.1175/1520-0434(1993)008<0281:WIAGFA>2.0.CO;2).
- 903
- 904 Murphy, A.H., 1997. Forecast verification, in: Katz, R.W., Murphy, A.H. (Eds.), Economic Value of Weather and Climate Forecasts. Cambridge
905 University Press, pp. 19 – 74. doi:<https://doi.org/10.1017/CB09780511608278.003>.
- 906
- 907 Murphy, A.H., Brown, B.G., Chen, Y.S., 1989. Diagnostic verification of temperature forecasts. Weather and Forecasting 4, 485–501. doi:[https://doi.org/10.1175/1520-0434\(1989\)004<0485:DVOTF>2.0.CO;2](https://doi.org/10.1175/1520-0434(1989)004<0485:DVOTF>2.0.CO;2).
- 908
- 909 Murphy, A.H., Winkler, R.L., 1971. forecasters and probability forecasts: some current problems,. Bulletin of the American Meteorological
910 Society 52, 239 – 248. doi:[https://doi.org/10.1175/1520-0477\(1971\)052<0239:FAPFSC>2.0.CO;2](https://doi.org/10.1175/1520-0477(1971)052<0239:FAPFSC>2.0.CO;2).
- 911
- 912 Murphy, A.H., Winkler, R.L., 1987. A general framework for forecast verification. Monthly Weather Review 115, 1330 – 1338. doi:[https://doi.org/10.1175/1520-0493\(1987\)115<1330:AGFFV>2.0.CO;2](https://doi.org/10.1175/1520-0493(1987)115<1330:AGFFV>2.0.CO;2).
- 913
- 914 Pedro, H.T.C., Coimbra, C.F.M., 2015. Short-term irradiance forecastability for various solar micro-climates. Solar Energy 122, 587
915 – 602. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X15005162>, doi:<https://doi.org/10.1016/j.solener.2015.09.031>.
- 916
- 917 Perez, R., Ineichen, P., Seals, R., Michalsky, J., Stewart, R., 1990. Modeling daylight availability and irradiance components from direct and
918 global irradiance. Solar Energy 44, 271 – 289. URL: <http://www.sciencedirect.com/science/article/pii/0038092X9090055H>,
919 doi:[https://doi.org/10.1016/0038-092X\(90\)90055-H](https://doi.org/10.1016/0038-092X(90)90055-H).
- 920
- 921 Perez, R., Lorenz, E., Pelland, S., Beauharnois, M., Knowe, G.V., Hemker, K., Heinemann, D., Remund, J., Müller, S.C., Traunmüller, W.,
922 Steinmauer, G., Pozo, D., Ruiz-Arias, J.A., Lara-Fanego, V., Ramirez-Santigosa, L., Gaston-Romero, M., Pomares, L.M., 2013. Comparison
923 of numerical weather prediction solar irradiance forecasts in the US, Canada and Europe. Solar Energy 94, 305 – 326. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X13001886>, doi:<https://doi.org/10.1016/j.solener.2013.05.005>.
- 924
- 925 Perez, R., Schlemmer, J., Kankiewicz, A., Dise, J., Tadese, A., Hoff, T., 2017. Detecting calibration drift at ground truth stations a demonstration
926 of satellite irradiance models' accuracy, in: 2017 IEEE 44th Photovoltaic Specialist Conference (PVSC), pp. 1104–1109. doi:<https://doi.org/10.1109/PVSC.2017.8366469>.
- 927
- 928 Perpiñán, O., Marcos, J., Lorenzo, E., 2013. Electrical power fluctuations in a network of DC/AC inverters in a large PV plant: Relationship
929 between correlation, distance and time scale. Solar Energy 88, 227 – 241. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X12004197>, doi:<https://doi.org/10.1016/j.solener.2012.12.004>.
- 930
- 931 Polo, J., Martín-Pomares, L., Gueymard, C.A., Balenzategui, J.L., Fabero, F., Silva, J.P., 2019. Fundamentals: Quantities, definitions, and units,

917 in: Polo, J., Martín-Pomares, L., Sanfilippo, A. (Eds.), *Solar Resources Mapping: Fundamentals and Applications*. Springer International
918 Publishing, Cham, pp. 1 – 14. doi:https://doi.org/10.1007/978-3-319-97484-2_1.

919 Ren, Y., Suganthan, P., Srikanth, N., 2015. Ensemble methods for wind and solar power forecasting—A state-of-the-art review. *Renewable
920 and Sustainable Energy Reviews* 50, 82 – 91. URL: <http://www.sciencedirect.com/science/article/pii/S1364032115003512>,
921 doi:<https://doi.org/10.1016/j.rser.2015.04.081>.

922 Reno, M.J., Hansen, C.W., 2016. Global horizontal irradiance clear sky models: Implementation and analysis. *Renewable Energy* 90, 520 – 531.
923 URL: <http://www.sciencedirect.com/science/article/pii/S0038092X15002200>, doi:[https://doi.org/10.1016/j.renene.
924 2015.12.031](https://doi.org/10.1016/j.renene.2015.12.031).

925 Ruiz-Arias, J.A., Gueymard, C.A., 2018a. A multi-model benchmarking of direct and global clear-sky solar irradiance predictions at arid sites using
926 a reference physical radiative transfer model. *Solar Energy* 171, 447 – 465. URL: [http://www.sciencedirect.com/science/article/
927 pii/S0038092X18305991](http://www.sciencedirect.com/science/article/pii/S0038092X18305991), doi:<https://doi.org/10.1016/j.solener.2018.06.048>.

928 Ruiz-Arias, J.A., Gueymard, C.A., 2018b. Worldwide inter-comparison of clear-sky solar radiation models: Consensus-based review of direct
929 and global irradiance components simulated at the earth surface. *Solar Energy* 168, 10 – 29. URL: [http://www.sciencedirect.com/
930 science/article/pii/S0038092X18301257](http://www.sciencedirect.com/science/article/pii/S0038092X18301257), doi:<https://doi.org/10.1016/j.solener.2018.02.008>. advances in Solar Resource
931 Assessment and Forecasting.

932 Schilling, R.L., 2017. *Measures, integrals and martingales*. Cambridge University Press.

933 Sengupta, M., Habte, A., Kurtz, S., Dobos, A., Wilbert, S., Lorenz, E., Stoffel, T., Renné, D., Gueymard, C.A., Myers, D., et al., 2015. Best
934 practices handbook for the collection and use of solar resource data for solar energy applications. Technical Report NREL/TP-5D00-63112.
935 National Renewable Energy Laboratory.

936 Sun, X., Bright, J.M., Gueymard, C.A., Acord, B., Wang, P., Engerer, N.A., 2019. Worldwide performance assessment of 75 global clear-sky
937 irradiance models using principal component analysis. *Renewable and Sustainable Energy Reviews* 111, 550 – 570. URL: [http://www.
938 sciencedirect.com/science/article/pii/S1364032119302187](http://www.sciencedirect.com/science/article/pii/S1364032119302187), doi:<https://doi.org/10.1016/j.rser.2019.04.006>.

939 Tashman, L.J., 2000. Out-of-sample tests of forecasting accuracy: an analysis and review. *International Journal of Forecasting* 16,
940 437 – 450. URL: <http://www.sciencedirect.com/science/article/pii/S0169207000000650>, doi:[https://doi.org/10.1016/
941 S0169-2070\(00\)00065-0](https://doi.org/10.1016/S0169-2070(00)00065-0). the M3- Competition.

942 Urraca, R., Gracia-Amillo, A.M., Huld, T., de Pison, F.J.M., Trentmann, J., Lindfors, A.V., Riihelä, A., Sanz-Garcia, A., 2017. Quality control of
943 global solar radiation data with satellite-based products. *Solar Energy* 158, 49 – 62.

944 Urraca, R., Huld, T., Lindfors, A.V., Riihelä, A., de Pison, F.J.M., Sanz-Garcia, A., 2018. Quantifying the amplified bias of PV system simulations
945 due to uncertainties in solar radiation estimates. *Solar Energy* 176, 663 – 677. URL: [http://www.sciencedirect.com/science/article/
946 pii/S0038092X18310442](http://www.sciencedirect.com/science/article/pii/S0038092X18310442), doi:<https://doi.org/10.1016/j.solener.2018.10.065>.

947 Vallance, L., Charbonnier, B., Paul, N., Dubost, S., Blanc, P., 2017. Towards a standardized procedure to assess solar forecast accuracy: A
948 new ramp and time alignment metric. *Solar Energy* 150, 408 – 422. URL: [http://www.sciencedirect.com/science/article/pii/
949 S0038092X17303687](http://www.sciencedirect.com/science/article/pii/S0038092X17303687), doi:<https://doi.org/10.1016/j.solener.2017.04.064>.

950 Voyant, C., Notton, G., 2018. Solar irradiation nowcasting by stochastic persistence: A new parsimonious, simple and efficient forecasting
951 tool. *Renewable and Sustainable Energy Reviews* 92, 343 – 352. URL: [http://www.sciencedirect.com/science/article/pii/
952 S1364032118303344](http://www.sciencedirect.com/science/article/pii/S1364032118303344), doi:<https://doi.org/10.1016/j.rser.2018.04.116>.

953 Voyant, C., Notton, G., Kalogirou, S., Nivet, M.L., Paoli, C., Motte, F., Fouilloy, A., 2017. Machine learning methods for solar radiation forecasting:
954 A review. *Renewable Energy* 105, 569 – 582. URL: <http://www.sciencedirect.com/science/article/pii/S0960148116311648>,
955 doi:<https://doi.org/10.1016/j.renene.2016.12.095>.

956 Wasserman, L., 2013. *All of statistics: a concise course in statistical inference*. Springer Science & Business Media.

957 Willmott, C.J., Matsuura, K., 2005. Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average
958 model performance. *Climate Research* 30, 79 – 82. URL: <https://www.int-res.com/abstracts/cr/v30/n1/p79-82/>, doi:<https://doi.org/10.3354/cr030079>.

960 Yang, D., 2016. Solar radiation on inclined surfaces: Corrections and benchmarks. *Solar Energy* 136, 288 – 302. URL: [http://www.
961 sciencedirect.com/science/article/pii/S0038092X16302432](http://www.sciencedirect.com/science/article/pii/S0038092X16302432), doi:<https://doi.org/10.1016/j.solener.2016.06.062>.

962 Yang, D., 2018. A correct validation of the National Solar Radiation Data Base (NSRDB). *Renewable and Sustainable Energy Reviews* 97,
963 152 – 155. URL: <http://www.sciencedirect.com/science/article/pii/S1364032118306087>, doi:[https://doi.org/10.1016/
964 j.rser.2018.08.023](https://doi.org/10.1016/j.rser.2018.08.023).

965 Yang, D., 2019a. A guideline to solar forecasting research practice: Reproducible, operational, probabilistic or physically-based, ensemble, and
966 skill (ROPES). *Journal of Renewable and Sustainable Energy* 11, 022701. doi:<https://doi.org/10.1063/1.5087462>.

967 Yang, D., 2019b. Making reference solar forecasts with climatology, persistence, and their optimal convex combination. *Solar Energy* 193,
968 981 – 985. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X19309880>, doi:[https://doi.org/10.1016/
969 j.solener.2019.10.006](https://doi.org/10.1016/j.solener.2019.10.006).

970 Yang, D., 2019c. Standard of reference in operational day-ahead deterministic solar forecasting. *Journal of Renewable and Sustainable Energy* 11,
971 053702. doi:<https://doi.org/10.1063/1.5114985>.

972 Yang, D., Boland, J., 2019. Satellite-augmented diffuse solar radiation separation models. *Journal of Renewable and Sustainable Energy* 11,
973 023705. doi:<https://doi.org/10.1063/1.5087463>.

974 Yang, D., Kleissl, J., Gueymard, C.A., Pedro, H.T.C., Coimbra, C.F.M., 2018. History and trends in solar irradiance and PV power forecasting:
975 A preliminary assessment and review using text mining. *Solar Energy* 168, 60–101. URL: [http://www.sciencedirect.com/science/
976 article/pii/S0038092X17310022](http://www.sciencedirect.com/science/article/pii/S0038092X17310022), doi:<https://doi.org/10.1016/j.solener.2017.11.023>. Advances in Solar Resource Assess-
977 ment and Forecasting.

978 Yang, D., Perez, R., 2019. Can we gauge forecasts using satellite-derived solar irradiance? *Journal of Renewable and Sustainable Energy* 11,
979 023704. doi:<https://doi.org/10.1063/1.5087588>.

980 Yang, D., Quan, H., Disfani, V.R., Liu, L., 2017. Reconciling solar forecasts: Geographical hierarchy. *Solar Energy* 146, 276
981 – 286. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17301020>, doi:<https://doi.org/10.1016/j.rser.2017.04.006>.

982 [solener.2017.02.010](https://doi.org/10.1016/j.solener.2017.02.010).

983 Yang, D., Sharma, V., Ye, Z., Lim, L.I., Zhao, L., Aryaputera, A.W., 2015. Forecasting of global horizontal irradiance by exponential smoothing,
984 using decompositions. *Energy* 81, 111 – 119. URL: <http://www.sciencedirect.com/science/article/pii/S0360544214013528>,
985 doi:<https://doi.org/10.1016/j.energy.2014.11.082>.

986 Zhang, J., Florita, A., Hodge, B.M., Lu, S., Hamann, H.F., Banunarayanan, V., Brockway, A.M., 2015. A suite of metrics for assessing the
987 performance of solar power forecasting. *Solar Energy* 111, 157 – 175. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0038092X14005027)
988 [S0038092X14005027](http://www.sciencedirect.com/science/article/pii/S0038092X14005027), doi:<https://doi.org/10.1016/j.solener.2014.10.016>.