

1 Data-Driven Model for Solar Irradiation Based on Satellite 2 Observations

3 Ilias Bilonis^a, Emil M. Constantinescu^a, Mihai Anitescu^a

4 ^a*Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439.*
5 *Emails: {ebilonis,anitescu,emconsta}@mcs.anl.gov*

6 Abstract

We construct a data-driven model for solar irradiation based on satellite observations. The model yields probabilistic estimates of the irradiation field every thirty minutes starting from two consecutive satellite measurements. The probabilistic nature of the model captures prediction uncertainties and can therefore be used by solar energy producers to quantify the operation risks. The model is simple to implement and can make predictions in realtime with minimal computational resources. To deal with the high-dimensionality of the satellite data, we construct a reduced representation using factor analysis. Then, we model the dynamics of the reduced representation as a discrete (30-minute interval) dynamical system. In order to convey information about the movement of the irradiation field, the dynamical system has a two-step delay. The dynamics are represented in a nonlinear, nonparameteric way by a recursive Gaussian process. The predictions of the model are compared with observed satellite data as well as with a similar model that uses only ground observations at the prediction site. We conclude that using satellite data in an area including the prediction site significantly improves the prediction compared with models using only ground observation site data.

7 *Keywords:*

8 insolation, irradiance, Bayesian, recursive Gaussian process, factor analysis, dimensionality
9 reduction, dynamical system

11 1. Introduction

12 Solar irradiation is the amount of power per square meter that reaches the Earth from the
13 Sun. In solar energy applications, part of the solar irradiation can be converted to electricity.
14 In contrast to conventional power sources such as coal or gas, solar irradiation is volatile and
15 uncontrollable by the user. The two most important factors of the solar irradiation variability are
16 the movement of the Sun and weather fluctuations. The former can be captured mathematically
17 to great accuracy, because it is a deterministic effect. The latter is a chaotic effect, and hence the
18 main cause for the difficulties associated with forecasting solar irradiation.

19 All the stages of a solar-power conversion project need to take into account the risks asso-
20 ciated with solar irradiation. For the feasibility and design phases of the project, historical data
21 can be employed to quantify these risks. The risks associated with the operation phase, how-
22 ever, require the ability to make short-term predictions (from 1 to 8 hours ahead) of the solar
23 irradiation.

24 The most widely used solar irradiation forecasting methodologies are those that rely only on
25 pointwise ground measurements of the solar irradiation. The reason is that ground measurements

26 are readily available at any solar energy production plant. Mathematically, these techniques fall
27 into the category of single-value time series analysis. Time series analysis methods include
28 the autoregressive integrated moving average (ARIMA) processes [1] (see [2] for a first-order
29 autoregressive model (AR(1)) and [3, Ch. 15.2.2.] for ARIMA examples), and artificial neural
30 networks (ANNs) [4]. The accuracy of the predictions of these models degrades rapidly as the
31 forecasting window is increased. This result is expected because the weather fluctuations exhibit
32 a nonlocal behavior; see [3, Ch. 15] for a comprehensive review.

33 More accurate forecasts can be achieved only if nonlocal data are taken into account. For
34 short-term (0-30 minutes) forecasts, a promising approach is to use total sky imager technolo-
35 gies [5]. One takes pictures of the sky from a particular site, extracts information about the
36 clouds, constructs the cloud motion vectors (CMVs) and moves the image forward in time. Us-
37 ing geometrical arguments and semi-empirical models, one recovers the solar irradiation from
38 the cloud information at that later time. The time frame for which this approach is useful depends
39 on the velocity of the clouds.

40 Longer forecasts (hours to days) are feasible if satellite data are used. For forecasts ranging
41 from 30 minutes to 6 hours ahead, the data-driven technique introduced in [6] may be used. As a
42 first step, the semi-empirical heliostat method of [7] is used to extract cloud structure information
43 from the satellite images. Then, as in the sky imager-based techniques, two consecutive images
44 are compared in order to construct the CMVs. The cloud information is moved forward in time
45 using the CMVs and goes through a final smoothing phase. Solar irradiation is recovered by
46 again employing the heliostat method. Numerical Weather Prediction (NWP) may be used for
47 forecasts of up to 6 days or longer. For example, in [3, Ch. 10] the sky-cover fraction of the U.S.
48 National Digital Forecast Database is coupled with semi-empirical models to produce long-term
49 forecasts of solar irradiation.

50 The main disadvantage of most these aforementioned techniques is that they are difficult to
51 use, and, sometimes, unable to quantify the uncertainty in their predictions. Given the current
52 evolution of decision systems for energy toward incorporating stochastic representations [8], this
53 may be a serious shortcoming. It is hard to see how to consistently add an uncertainty model
54 to the heliostat approach. NWP models can in principle be modified to support an ensemble-
55 based approach to uncertainty, but at a significant computational cost that requires a dedicated
56 supercomputer [9]. Among the methods described, only the ARIMA-based methods can provide
57 error bars for the predictions with a small or moderate effort. Yet, it is exactly these error bars
58 that help quantify the potential risks and allow the stakeholders to properly price them.

59 These considerations have motivated us to develop a fully stochastic model that can quantify
60 forecast uncertainties. In addition, aiming for a model that is convenient even for lean operations,
61 we propose a satellite-based model that is considerably more robust than existing ones and can
62 produce predictions in realtime with minimal computational resources.

63 Our model building philosophy and paper can be summarized as follows. As a first step, the
64 Sun's movement effect on the satellite observed solar irradiation field (Sec. 2.1) is removed by
65 dividing it with a clear sky model (Sec. 2.2) to get the clear sky index field. Our goal is to use
66 consecutive observations of the clear sky index field in order to learn its dynamics. Because of its
67 high-dimensional nature, we construct a reduced-dimensionality representation of it (Sec. 2.3).
68 To learn the dynamics of this low-dimensional representation, we use a nonlinear, nonparametric
69 technique known as recursive Gaussian process (Sec. 2.4). Having constructed the dynamics
70 of the reduced space, forecasts can be performed for an arbitrary number of time steps ahead
71 (Sec. 2.5). Our recursive Gaussian process is similar in concept to the ANN used in [4]. However,
72 our model is Bayesian, a key feature that enables us to make not only best estimates but also

73 probabilistic forecasts. We then present (Sec. 3) our numerical results for an 8-hours-ahead
74 forecast, and we compare them pointwise with those obtained by a recursive Gaussian process
75 model based only on ground observations. We observe that using satellite data significantly
76 reduces the forecast uncertainties and improves the forecast itself. We attribute this improvement
77 to the nonlocal information carried by the satellite images and to the space-time correlation
78 between the solar irradiation at the prediction site and at neighboring sites.

79 2. Methodology

80 Throughout this work, we denote the solar irradiation field by $I(\phi, \lambda, t)$, where ϕ, λ and t are
81 the latitude, longitude, and time, respectively. The units of $I(\phi, \lambda, t)$ are power per square meter.

82 2.1. Observations of Solar Irradiation

83 The solar irradiation field can be observed almost instantaneously by processing satellite im-
84 ages. For simplicity, we restrict our attention to data coming from the continental United States
85 (CONUS) scan of the GOES-13 satellite. The CONUS scan takes place almost every 30 min-
86 utes and has a resolution of 1 km. The solar irradiation field is constructed by processing the
87 raw data collected by the Advanced Very High Resolution Radiometer (AVHRR) via the algo-
88 rithms described in [10]. This product is reported as part of the Clouds from AVHRR Extended
89 (CLAVR-x, <https://cimss.ssec.wisc.edu/clavr/>) data and is available in realtime from
90 the University of Wisconsin (ftp://ftp.ssec.wisc.edu/clavr/goes_east/processed/).
91 The CLAVR-x data are available on a grid of latitudes and longitudes described by ϕ_{ij} and λ_{ij} for
92 $i = 1, \dots, P_\phi, j = 1, \dots, P_\lambda$, where P_ϕ and P_λ denote the number of pixels on each dimension.
93 That is, at time t we observe a matrix $\mathbf{I}(t) = (I_{ij}(t))$ of size $P_\phi \times P_\lambda$:

$$I_{ij}(t) = I(\phi_{ij}, \lambda_{ij}, t). \quad (1)$$

94 The dimensions of the observed matrix for the CONUS scan are $P_\phi = 1,900, P_\lambda = 3,100$. For
95 computational reasons, we work with a submatrix of the CONUS scan of dimensions $P_\phi = P_\lambda =$
96 400 that is centered on a site of interest. We refer to this submatrix as the *patch*.

97 The time resolution of the data is irregular. In particular, we have observations every thirty
98 minutes, with the exception of 03:00, 06:00, 09:00, 12:00, 15:00, 15:30, 18:00, 21:00, and 00:00
99 UTC daily as well as 17:00 UTC every Wednesday. In addition, early morning and late after-
100 noon measurements are not useful, because of the very low irradiation values that are observed.
101 Therefore, we drop from our dataset observations that take place during nighttime (i.e., the pe-
102 riod between two hours before sunset up to two hours after sunrise the next day, where sunset
103 and sunrise times are calculated for an observer standing at the center of the patch). We denote
104 the times for which we actually have observations by $t_k, k = 1, \dots, K_t$. Then, our observed dataset
105 consists of

$$\mathcal{D}_I = \{\mathbf{I}(t_k) : k = 1, \dots, K_t\}. \quad (2)$$

106 2.2. Clear Sky Model and the Clearness Index

107 The presence of daily and seasonal trends makes the direct modeling of $I(\phi, \lambda, t)$ problematic.
108 On the one hand, the daily part is characterized by (1) no variation over night (the solar irradiation
109 field is equal to zero), (2) increasing values and variance until noon, and (3) decreasing values
110 and variance from noon till sunset. On the other hand, the seasonal trend is characterized by (1)

low values during winter and high values during summer, and (2) high variance during winter and low variance during summer. The change in the mean values of the field is directly chained to the solar zenith angle, which varies daily as well as seasonally. The higher variance during winter days is attributed to the rapidly changing weather conditions. To partially address these trends, we model the *clearness index* instead of the solar irradiation field. In order to properly define the clearness index, the concept of a *clear sky model* is required.

A clear sky model approximates the expected solar irradiation when no clouds are present. Typically, it depends on the extraterrestrial irradiation ((W/m^2) at the top of Earth's surface), the solar zenith angle at a particular Earth site, the elevation of the site above sea level, the composition of atmospheric gases such as water vapor and ozone content, and the atmospheric aerosol content (see [3, 2.3]). Many clear sky models may be used (e.g., [11, 12]). However, given the variability in the total solar irradiation due to the cloud variability, very high accuracy of the clear sky model is not extremely important here. We thus opt for Ineichen's model [12], which is a surrogate of the more accurate Solis model [13]. The Ineichen model depends on the extraterrestrial irradiation, the solar zenith, the elevation of the site above sea level, the atmospheric water vapor content, and the aerosol optical depth at 700 nm. We denote the global clear sky irradiation of Ineichen's model by $I_{\text{cls}}(\phi, \lambda, t)$. Details on the various model parameters are given in Appendix A.

We can now define the clearness index field as the ratio of the solar irradiation field to the clear sky model:

$$C(\phi, \lambda, t) = \frac{I(\phi, \lambda, t)}{I_{\text{cls}}(\phi, \lambda, t)}. \quad (3)$$

The observed data of Eq. (1) may be expressed in terms of the clearness index by

$$\mathcal{D}_C = \{\mathbf{C}(t_k) : k = 1, \dots, K_t\}, \quad (4)$$

where $\mathbf{C}(t)$ is the $P_\phi \times P_\lambda$ -matrix defined as

$$C_{ij}(t) = \frac{I_{ij}(t)}{I_{\text{cls},ij}(t)} = \frac{I(\phi_{ij}, \lambda_{ij}, t)}{I_{\text{cls}}(\phi_{ij}, \lambda_{ij}, t)},$$

for $i = 1, \dots, P_\phi, j = 1, \dots, P_\lambda$.

Our goal is to model the evolution of the clearness index field $C(\phi, \lambda, t)$ based on the observations included in Eq. (4). In reality, $C(\phi, \lambda, t)$ is a continuous spatio-temporal process. However, we will be focusing on a discrete version of the problem. Namely, we will be modeling the evolution of the matrix $\mathbf{C}(t)$ over thirty minute intervals. Since the clear sky model may be evaluated anywhere at will, the solar irradiation field can be recovered in a trivial way from Eq. (3).

We note that some trends do persist even after switching to the clearness index. In particular, the clearness index field exhibits a seasonal trend in its variance similar to that of the solar irradiation field. However, this remaining trend is inconsequential, because our goal is to make short-term forecasts.

2.3. Dealing with the Curse of Dimensionality

The mathematical problem we are facing is the data-driven determination of the dynamics of an extremely high-dimensional system. The very small patch of data we are considering has 160,000 ((400×400)) dimensions. Attempting to learn the dynamics of the clearness index directly is not tractable. Clearly, we must resort to some kind of *dimensionality reduction technique*.

149 The observed data are in the high-dimensional space $\mathbb{R}^{P_\phi \times P_\lambda}$. However, we expect them to
 150 be constrained on a low-dimensional manifold embedded in this high-dimensional space. This
 151 expectation is based on the spatial correlations of the field that are induced by physical laws in-
 152 volving cloud transportation, nucleation, and annihilation. In the hypothetical scenario of a truly
 153 $P_\phi \times P_\lambda$ -dimensional manifold, there would be no spatial correlations. But such a scenario is not
 154 supported by the data: displaying successive maps of the solar irradiation field shows evolution
 155 of persistent moderate scale patterns, which is consistent with nontrivial spatial correlation.

156 To be precise, let $\mathcal{M}_h \subset \mathbb{R}^{P_\phi \times P_\lambda}$ denote the manifold in which the clearness index lives.
 157 Mathematically, it is defined as

$$\mathcal{M}_h = \{\mathbf{C}(t) : \text{for all } t\}. \quad (5)$$

158 We expect the true dimensionality of this manifold to be less than $P_\phi \times P_\lambda$. That is, we expect
 159 that it can be parametrized with fewer than $P_\phi \times P_\lambda$ variables.

160 A dimensionality reduction technique should provide two maps: a *reduction* map, and a
 161 *reconstruction* map from and to \mathcal{M}_h , respectively. A reduction map $\mathcal{R} : \mathcal{M}_h \rightarrow \mathbb{R}^R$ projects
 162 the original data to the low-dimensional space \mathbb{R}^R with $R \ll P_\phi \times P_\lambda$. A *reconstruction map*
 163 $\mathcal{C} : \mathbb{R}^R \rightarrow \mathcal{M}_h$ maps the low-dimensional space back to the original one. A *good* pair of
 164 reduction-reconstruction maps $(\mathcal{R}, \mathcal{C})$ should have an R as small as possible and also satisfy the
 165 following property,

$$\mathcal{C}(\mathcal{R}(\mathbf{C}(t))) \approx \mathbf{C}(t), \quad (6)$$

166 with “ \approx ” being interpreted as “close” in some appropriate norm. In other words, the reconstruc-
 167 tion map should approximately be the inverse of the reduction map. In reality, we always expect
 168 some information loss along the way. After finding a good $(\mathcal{R}, \mathcal{C})$ pair, our goal will be to model
 169 the dynamics of the R -dimensional projections of the original data. This is a much easier task
 170 than dealing directly with $\mathbf{C}(t)$. To clarify this point, let us denote by $\mathbf{x}(t) \in \mathbb{R}^R$ the *reduced*
 171 variables:

$$\mathbf{x}(t) = \mathcal{R}(\mathbf{C}(t)). \quad (7)$$

172 Our goal is to capture the dynamics of $\mathbf{x}(t)$. The original dynamics can be recovered by passing
 173 $\mathbf{x}(t)$ through the reconstruction map. The observations we have at hand for this are induced by
 174 Eq. (4) through the reduction map

$$\mathcal{D}_x = \{\mathbf{x}(t_k) = \mathcal{R}(\mathbf{C}(t_k)) : k = 1, \dots, K_t\}. \quad (8)$$

175 This is the topic of Sec. 2.4. For the moment, we focus on finding $(\mathcal{R}, \mathcal{C})$.

176 The only information available for determining the pair $(\mathcal{R}, \mathcal{C})$ is the observed clearness index
 177 dataset \mathcal{D}_C given in Eq. (4). The dimensionality reduction technique we will use is the factor
 178 analysis (FA)[14, 15, 16]. FA can be thought of as a generalization of probabilistic principal
 179 component analysis (pPCA) [17, 18], a probabilistic interpretation of the celebrated principal
 180 component analysis (PCA) (see [19, 20] and for a more recent reference [21, Ch. 12.1]). PCA
 181 works under the assumption that the data manifold \mathcal{M}_h is linear. In other words, one may think
 182 of it as an attempt to approximate \mathcal{M}_h as being embedded in a low-dimensional affine space. In
 183 addition, pPCA and FA attempt to characterize the uncertainty of the reduction/reconstruction
 184 operations. Both fall into the category of latent linear Gaussian models; but pPCA uses the same
 185 variance to characterize the reconstruction uncertainty, whereas FA uses a different variance for
 186 each distinct feature.

187 FA works with vectors. Therefore, all quantities must be vectorized before proceeding fur-
 188 ther. For notational convenience, let $\text{vec}_{m,n} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{mn}$ be the *vectorization* operator. The
 189 action of the vectorization operator on a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ transforms it into a vector $\mathbf{a} \in \mathbb{R}^{mn}$, as
 190 follows:

$$\mathbf{a} = \text{vec}_{m,n}(\mathbf{A}) := (A_{11}, \dots, A_{1n}, \dots, A_{m1}, \dots, A_{mn}). \quad (9)$$

191 In addition, we need the inverse transform $\text{vec}_{m,n}^{-1}$, which transforms a vector $\mathbf{a} \in \mathbb{R}^{mn}$ into a
 192 matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$:

$$\mathbf{A} = \text{vec}_{m,n}^{-1}(\mathbf{a}) := \begin{pmatrix} a_1 & \dots & a_n \\ \vdots & \ddots & \vdots \\ a_{m(n-1)+1} & \dots & a_{mn} \end{pmatrix}. \quad (10)$$

193 Since we are interested in vectorizing $\mathbf{C}(t) \in \mathbb{R}^{P_\phi \times P_\lambda}$, we need to use $\text{vec}_{P_\phi \times P_\lambda}$. For brevity, we
 194 define the following:

$$\text{vec} := \text{vec}_{P_\phi \times P_\lambda} \text{ and } \text{vec}^{-1} := \text{vec}_{P_\phi \times P_\lambda}^{-1}. \quad (11)$$

195 Our problem now is to find a reduced representation of the K_t ($P_\phi P_\lambda$)-dimensional observations:

$$\mathcal{D}_c := \{\mathbf{c}(t_k) = \text{vec}(\mathbf{C}(t_k)) : k = 1, \dots, K_t\}. \quad (12)$$

197 In our discussion of FA we follow closely the work of Bishop [21, Ch. 12.2.4]. Let $\mathbf{x} \in \mathbb{R}^R$
 198 be the latent reduced representation of an observed clearness index vector $\mathbf{c} \in \mathbb{R}^{P_\phi P_\lambda}$. In FA, we
 199 assume that the observed clearness vector \mathbf{c} is generated from \mathbf{x} as follows,

$$p(\mathbf{c}|\mathbf{x}, \mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}) = \mathcal{N}(\mathbf{c}|\mathbf{W}\mathbf{x} + \boldsymbol{\mu}, \boldsymbol{\Psi}), \quad (13)$$

200 where $\mathbf{W} \in \mathbb{R}^{(P_\phi P_\lambda) \times R}$ and $\boldsymbol{\mu} \in \mathbb{R}^{P_\phi P_\lambda}$ are the *reconstruction matrix*, and the *mean vector*, respec-
 201 tively, and $\boldsymbol{\Psi} \in \mathbb{R}^{(P_\phi P_\lambda) \times (P_\phi P_\lambda)}$ is a diagonal matrix containing the *reconstruction variance* of each
 202 component of \mathbf{c} (each component is known as a *feature*). Equation (13) defines a probabilistic
 203 version of the reconstruction map \mathcal{C} .

204 To close the model, one must say how \mathbf{x} would be generated if \mathbf{c} was not observed. In FA,
 205 one assumes that \mathbf{x} is generated from a unit variance Gaussian,

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{I}), \quad (14)$$

206 where \mathbf{I} is the unit matrix.

207 The reduction map \mathcal{R} is also defined probabilistically. One has to compute the posterior of \mathbf{x}
 208 conditioned on \mathbf{c} and all the parameters $(\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi})$. This is again a Gaussian of the form

$$p(\mathbf{x}|\mathbf{c}, \mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}) = \mathcal{N}(\mathbf{x}|\mathbf{G}\mathbf{W}^T\boldsymbol{\Psi}^{-1}(\mathbf{c} - \boldsymbol{\mu}), \mathbf{G}), \quad (15)$$

209 where

$$\mathbf{G} = (\mathbf{I} + \mathbf{W}^T\boldsymbol{\Psi}^{-1}\mathbf{W})^{-1}. \quad (16)$$

210 FA can be trained efficiently by maximizing the likelihood of \mathcal{D}_c of Eq. (12) using the
 211 expectation-maximization (EM) algorithm [22]. In particular, combining Eqs. (13) and (14) and
 212 assuming independence of the vector \mathbf{c} conditional on the projection on principal components \mathbf{x} ,
 213 we need to solve the following maximization problem:

$$\mathbf{W}^*, \boldsymbol{\mu}^*, \boldsymbol{\Psi}^* = \arg \max_{\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}} p(\mathcal{D}_c|\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}), \quad (17)$$

214 where the likelihood term is given by

$$p(\mathcal{D}_c|\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}) = \prod_{k=1}^{K_c} p(\mathbf{c}(t_k)|\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}), \quad (18)$$

215 with

$$p(\mathbf{c}(t_k)|\mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}) = \int p(\mathbf{c}(t_k)|\mathbf{x}(t_k), \mathbf{W}, \boldsymbol{\mu}, \boldsymbol{\Psi}) p(\mathbf{x}(t_k)) d\mathbf{x}(t_k). \quad (19)$$

216 From this point on, for notational convenience we write

$$\boldsymbol{\phi}^* = (\mathbf{W}^*, \boldsymbol{\mu}^*, \boldsymbol{\Psi}^*) \quad (20)$$

217 to denote the solution of the maximization problem of Eq. (17). The projection operator is then
218 stated as

$$\mathcal{R}(\mathbf{C}) := \mathbf{G}^* \mathbf{W}^{*T} (\boldsymbol{\Psi}^*)^{-1} (\text{vec}(\mathbf{C}) - \boldsymbol{\mu}^*), \quad (21)$$

219 with \mathbf{G}^* as in Eq. (16) evaluated for $\mathbf{W} = \mathbf{W}^*$ and $\boldsymbol{\Psi} = \boldsymbol{\Psi}^*$.

220 Certainly the assumption of independence of \mathbf{c} conditional on \mathbf{x} is arguable, although the
221 problem would become much harder if we did not use it. On the other hand, should \mathbf{c} be en-
222 tirely contained in the space spanned by \mathbf{x} for all times t , the assumption would be satisfied.
223 The assumption is thus less restrictive the more accurate the PCA projection is. Therefore, we
224 anticipate the conditional independence being a good approximation for forecasting large-scale
225 behavior, which tends to be contained in the principal components.

226 We now come to the problem of determining the appropriate value for R (the dimension of
227 the reduced space). This could be attacked in a fully Bayesian manner by following the ideas
228 of [23]. In this work, however, we simply fix R to the maximum value we can afford, given the
229 constraints of our model. In particular, as R increases so does the number of parameters of the
230 reduced dynamics model (see Sec. 2.4). Because of the limited number of observations we have,
231 it does not make sense to select an R that is greater than 8. In particular, for $R = 8$ the proposed
232 model has 35 parameters (see Sec. 3) that need to be determined from about 250 observations.

233 2.4. Learning the Reduced Dynamics

234 The task of this section is to learn the dynamics of the reduced variables $\mathbf{x}(t)$ (see Eq. (7))
235 based on \mathcal{D}_x (see Eq. (8)). We use a model for non-linear time-series known as a *recursive*
236 or *dynamic* Gaussian process [24, 25]. Such models have been used extensively over the past
237 decade in diverse tasks including human motion modeling/tracking [26, 27] and nonlinear signal
238 processing [28]. For notational clarity, we assume that the time t counts units of 30 minutes; that
239 is, $t - 1$ means 30 minutes before t , $t - 2$ means 60 minutes before t , and so forth. We assume
240 that the evolution of $\mathbf{x}(t)$ obeys the following discrete dynamics:

$$\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t-1), \mathbf{x}(t-2)), \quad (22)$$

241 where $\mathbf{f} : \mathbb{R}^R \times \mathbb{R}^R \rightarrow \mathbb{R}^R$ is an unknown function to be determined from the observations \mathcal{D}_x
242 of Eq. (8). The reason we have assumed a dependence on $\mathbf{x}(t-2)$ in addition to $\mathbf{x}(t-1)$ is to
243 include some information about the velocity field that transports the clearness index. Simply
244 including $\mathbf{x}(t-1)$ would yield a Markovian model that would not be able to capture any transport
245 properties. One can envision models that take into account $\mathbf{x}(t-3)$, $\mathbf{x}(t-4)$, and so on. The rest
246 of the discussion can be extended to cover this generic case in a straightforward manner. The

247 computational cost of such an approach increases only moderately. However, the amount of data
 248 required in order to train the model does increase significantly mostly because of the need of
 249 consecutive measurements.

250 The idea of a recursive Gaussian process is to represent the function \mathbf{f} as a Gaussian process
 251 to be learned from the data. In turn, a Gaussian process may be thought as a nonparametric way
 252 of performing regression tasks [29]. We denote the input of \mathbf{f} collectively by $\mathbf{z}(t) \in \mathbb{R}^{2R}$, where

$$\mathbf{z}(t) = (\mathbf{x}(t-1), \mathbf{x}(t-2)). \quad (23)$$

253 Its output is simply $\mathbf{x}(t)$. The data we have available for learning \mathbf{f} are

$$\mathcal{D}_R = \{(\mathbf{z}(t_k), \mathbf{x}(t_k)) : t_k - t_{k-1} = 1, t_{k-1} - t_{k-2} = 1, k = 3, \dots, K_t\}.$$

254 That is, each of the $(\mathbf{z}(t), \mathbf{x}(t))$ pairs requires exactly three consecutive measurements of the clear-
 255 ness index to be available. For future reference, we assume that there are N observations in \mathcal{D}_R
 256 as follows:

$$\mathcal{D}_R = \{(\mathbf{z}^{(i)}, \mathbf{x}^{(i)}) : i = 1, \dots, N\}. \quad (24)$$

257 The core methodology for Gaussian processes aims at learning real functions, that is, func-
 258 tions with one output. In contrast, here we are dealing with the problem of learning a multioutput
 259 function. Despite the fact that there is a wealth of methods for learning multioutput functions
 260 [30, 31, 32, 33]; we chose a simple approach that treats each output dimension of \mathbf{f} , indepen-
 261 dently. In particular, we assume that each of the components f_r , $r = 1, \dots, R$ of \mathbf{f} is a Gaussian
 262 process representing the evolution of the r th principal component

$$x_r(t) = f_r(\mathbf{z}(t)). \quad (25)$$

263 Each of these functions can be learned from a fraction of the data contained in \mathcal{D}_R of Eq. (24).
 264 In particular, f_r can be learned from

$$\mathcal{D}_{R,r} = \{(\mathbf{z}^{(i)}, x_r^{(i)}) : i = 1, \dots, N\}, \quad (26)$$

265 where x_r is the r th component of \mathbf{x} .

266 At this point, we elaborate on how each of the functions f_r is constructed based on $\mathcal{D}_{R,r}$ of
 267 Eq. (26) using Gaussian process regression. For notational convenience, we drop the index r
 268 from any equation in this paragraph. For more details, the reader may consult [34]. Prior to
 269 seeing the data, we assume that the f_r is a draw from a zero mean Gaussian process,

$$f_r \sim \text{GP}(f_r | 0, k(\cdot, \cdot; \boldsymbol{\theta}_r)), \quad (27)$$

270 where $k : \mathbb{R}^{2R} \times \mathbb{R}^{2R} \rightarrow \mathbb{R}$ is a covariance function and $\boldsymbol{\theta}_r$ all its parameters. Notice that the same
 271 covariance function will be used for each r but its parameters $\boldsymbol{\theta}_r$ will be different. The particular
 272 form of the covariance functions used in the numerical examples is discussed in Sec. 3. The
 273 likelihood of the observed data is

$$p(\mathcal{D}_{R,r} | \boldsymbol{\theta}) = \mathcal{N}(\mathbf{X}_r | \mathbf{0}, \mathbf{A}_r), \quad (28)$$

274 where \mathbf{X}_r is the vector of all the x_r 's observed in $\mathcal{D}_{R,r}$ of Eq. (26),

$$\mathbf{X}_r = (x_r^{(1)}, \dots, x_r^{(N)}), \quad (29)$$

275 and $\mathbf{A}_r \in \mathbb{R}^{N \times N}$ is the covariance matrix evaluated at $\boldsymbol{\theta}_r$:

$$\mathbf{A}_{r,mm} = k(\mathbf{z}^{(n)}, \mathbf{z}^{(m)}; \boldsymbol{\theta}_r). \quad (30)$$

276 To train the model, we maximize the logarithm of the likelihood, Eq. (28), with respect to $\boldsymbol{\theta}_r$,

$$\boldsymbol{\theta}_r^* = \arg \max_{\boldsymbol{\theta}_r} \log p(\mathcal{D}_{R,r} | \boldsymbol{\theta}_r), \quad (31)$$

277 subject to any constraints we might have. Having found a point estimate $\boldsymbol{\theta}_r^*$ of $\boldsymbol{\theta}_r$, we have as the
278 *predictive distribution*

$$p(x_r(t) | \mathbf{z}(t), \mathcal{D}_{R,r}, \boldsymbol{\theta}_r^*) = \mathcal{N}(x_r(t) | m^*(\mathbf{z}(t)), \sigma^{*2}(\mathbf{z}(t))), \quad (32)$$

279 where m^* is the *predictive mean*,

$$m^*(\mathbf{z}(t)) = \mathbf{a}_r^*(\mathbf{z}(t)) \mathbf{A}_r^{*-1} \mathbf{X}_r, \quad (33)$$

280 and σ^{*2} is the *predictive variance*,

$$\sigma^{*2}(\mathbf{z}(t)) = k(\mathbf{z}(t), \mathbf{z}(t); \boldsymbol{\theta}_r^*) - \mathbf{a}_r^*(\mathbf{z}(t)) \mathbf{A}_r^{*-1} \mathbf{a}_r^*(\mathbf{z}(t))^T, \quad (34)$$

281 where \mathbf{A}_r^* is the covariance matrix of Eq. (30) evaluated at $\boldsymbol{\theta}_r^*$ and $\mathbf{a}_r^*(\mathbf{z}(t))$ is the *cross-covariance*:

$$\mathbf{a}_r^*(\mathbf{z}(t)) = \left(k(\mathbf{z}(t), \mathbf{z}^{(1)}; \boldsymbol{\theta}_r^*), \dots, k(\mathbf{z}(t), \mathbf{z}^{(N)}; \boldsymbol{\theta}_r^*) \right). \quad (35)$$

283 Combining the results for every r , we get the *multioutput predictive distribution*:

$$p(\mathbf{x}(t) | \mathbf{z}(t), \mathcal{D}_R, \boldsymbol{\theta}^*) = \prod_{r=1}^R p(x_r(t) | \mathbf{z}(t), \mathcal{D}_{R,r}, \boldsymbol{\theta}_r^*), \quad (36)$$

284 where $\boldsymbol{\theta}^* = \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_R\}$. Equation (36) characterizes in a probabilistic manner all the information
285 about the dynamics that we are able to get from the observations. It must be noted at this point,
286 that the many-step-ahead predictive distribution that is generated by recursive application of
287 Eq. (36) is non-Gaussian. This is due to the fact that the mean of Eq. (32) is nonlinear as a
288 function of $\mathbf{z}(t)$. In Sec. 2.5, we show how iteratively taking samples from Eq. (36) allows us to
289 make probabilistic forecasts about the evolution of the reduced dynamics.

290 2.5. Forecasting the Quantities of Interest

291 Now we come to the problem of forecasting the clearness index. Our predictions are going
292 to be probabilistic, and they will be based solely on Eq. (13), Eq. (15), and Eq. (36).

293 Assume that we have observed the clearness index at times t_0 and $t_0 - 1$; in other words,
294 we have observed the matrices $\mathbf{C}(t_0)$ and $\mathbf{C}(t_0 - 1)$, respectively. We would like to predict the
295 clearness index at future times $t = t_0 + 1, t_0 + 2, \dots$. First we need to project the observations to
296 the reduced components using Eq. (15),

$$\begin{aligned} \mathbf{x}(t_0) &\sim p(\mathbf{x}(t_0) | \text{vec}(\mathbf{C}(t_0)), \boldsymbol{\phi}^*), \\ \mathbf{x}(t_0 - 1) &\sim p(\mathbf{x}(t_0 - 1) | \text{vec}(\mathbf{C}(t_0 - 1)), \boldsymbol{\phi}^*), \end{aligned} \quad (37)$$

297 where $\boldsymbol{\phi}^*$ is defined in Eq. (20), or, in short notation,

$$\mathbf{z}(t_0) = (\mathbf{x}(t_0 - 1), \mathbf{x}(t_0 - 2)). \quad (38)$$

298 Note that this kind of probabilistic projection accounts for the uncertainty of the reduction map.

299 We can take sample paths from the second-order Markov chain defined by the predictive
 300 distribution of Eq. (36) to produce a sample path n steps long (see Alg. 1). The initial points
 301 $\mathbf{x}(t_0 - 1)$ and $\mathbf{x}(t_0)$ are assumed to be sampled as in Eq. (37). That is, we can produce sample
 302 paths that go arbitrarily far into the future, albeit in multiples of 30 minutes. Suppose that we have
 303 taken $S > 0$ such samples. Let us denote each of them by $\{\mathbf{x}^{(s)}(t_0 + i)\}_{i=1}^n, s = 1, \dots, S$. These
 304 can be transformed to sample paths of the clearness index matrix $\{\mathbf{C}^{(s)}(t_0 + i)\}_{i=1}^n, s = 1, \dots, S$
 305 using Eq. (13):

$$\begin{aligned} \mathbf{c}^{(s)}(t_0 + i) &\sim p(\mathbf{c}^{(s)} | \mathbf{x}^{(s)}(t_0 + i), \boldsymbol{\phi}^*), \\ \mathbf{C}^{(s)}(t_0 + i) &= \text{vec}^{-1}(\mathbf{c}^{(s)}(t_0 + i)). \end{aligned} \quad (39)$$

306 Equation 13 accounts for the uncertainty of the reconstruction map. If one wants sample paths
 307 $\{\mathbf{I}^{(s)}(t_0 + i)\}_{i=1}^n, s = 1, \dots, S$, of the solar irradiation, these can be trivially acquired by using
 308 Eq. (3),

$$\mathbf{I}^{(s)}(t_0 + i) = \mathbf{C}^{(s)}(t_0 + i) * \mathbf{I}_{\text{cls}}(t_0 + i), \quad (40)$$

309 where $\mathbf{I}_{\text{cls}}(t_0 + i)$ is the $P_\phi \times P_\lambda$ matrix containing the evaluation of the clear sky model at time
 310 $t_0 + i$ over the whole patch and “*” stands for *elementwise* multiplication of two matrices of the
 311 same dimensions.

Algorithm 1 Algorithm taking a single sample path from the second-order Markov chain defined by Eq. (36)

Require: Observed data: $\mathbf{x}(t_0), \mathbf{x}(t_0 - 1)$, number of steps ahead to predict: $n \geq 1$.

Ensure: $\{\mathbf{x}(t_0 + i)\}_{i=1}^n$ is a sample path of the second order Markov chain defined by Eq. (36).

for $i = 1$ to n **do**

Sample $\mathbf{x}(t_0 + i)$ from $p(\mathbf{x}(t_0 + i) | \mathbf{x}(t_0 + i - 1), \mathbf{x}(t_0 + i - 2), \mathcal{D}_R, \boldsymbol{\theta}^*)$ as defined in Eq. (36).

end for

312 3. Numerical Examples

313 As discussed in Sec. 2.1, the satellite observations of solar irradiation are obtained from the
 314 CLAVR-x project. We have been systematically storing the CLAVR-x data since October 10,
 315 2013. We will use all valid observations until February 25, 2014, to train our model. There
 316 are 822 such observations. The number is reduced because of two facts: (1) only observations
 317 that take place 2 hours after sunrise and 2 hours before sunset are retained, (2) there is a gap in
 318 observations during January 2014 because of technical problems faced by the CLAVR-x group.
 319 This constitutes our observed solar irradiation dataset \mathcal{D}_I of Eq. (2). From \mathcal{D}_I , the observed
 320 clearness index dataset \mathcal{D}_C of Eq. (4) is constructed as outlined in Sec. 2.2. The available data
 321 from February 26, 2014, to May 5, 2014, are used to test our predictions. In particular, we
 322 observe two consecutive measurements of the solar irradiation, we forecast the next 8 hours, and
 323 we compare the forecast to what was actually observed.

324 Because we would like to compare the performance of our model with one based on ground
 325 observations, we center our 400×400 pixel patch on Lamont, OK. For this particular location
 326 there exist excellent ground observations of solar irradiation that we can use for validation [35].
 327 The ground model to which we compare our results is again a recursive GP on the clearness

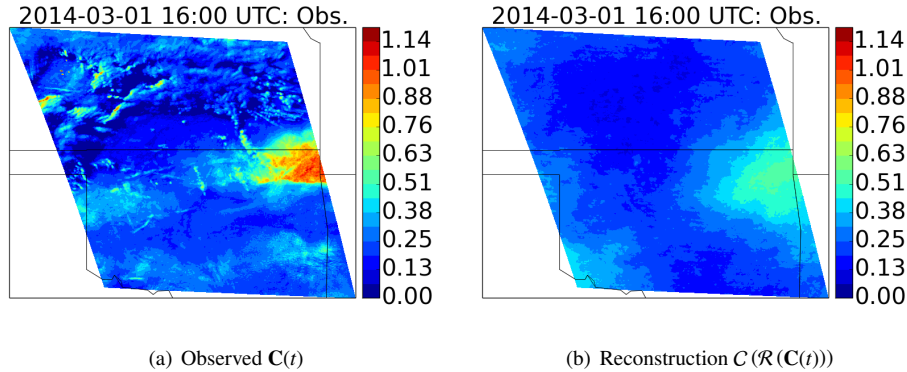


Figure 1: Sample of the clearness index on a 400×400 patch around Lamont, OK, on March 3, 2014, at 16:00 UTC (a) and its reconstruction with $R = 8$ components (b).

328 index observed at this site. The ground data used for training span the period from October 1,
 329 2013, to February 25, 2014.

330 The complete satellite data used in this work can be downloaded from: [http://www.mcs.
 331 anl.gov/~ebilionis/solar-irradiance/insolation_ok_400x400.h5](http://www.mcs.anl.gov/~ebilionis/solar-irradiance/insolation_ok_400x400.h5). In addition, we
 332 have created a website <http://www.mcs.anl.gov/~ebilionis/solar-irradiance/> that
 333 contains all 315 forecasts we could make between February 26, 2014 and May 5, 2014. Our
 334 model is implemented in Python and the code freely available at [https://bitbucket.org/
 335 ebilionis/solar-irradiation](https://bitbucket.org/ebilionis/solar-irradiation). The code contained therein, replicates all the figures of this
 336 paper as well as everything that can be found in the accompanying website.

337 The observed clearness index dataset \mathcal{D}_C is used to perform the dimensionality reduction
 338 task (see Sec. 2.3). In our numerical examples, we have experimented with $R = 4$ and $R = 8$
 339 components without observing any measurable performance gains by adding more complexity.
 340 For $R = 16$ components, the likelihood maximization encounters many local maxima because the
 341 number of parameters is considerably larger. In this case, either a global approach to maximizing
 342 the likelihood or some form of regularization, for example a fully Bayesian treatment with priors
 343 on the covariance parameters in (27), is required. We will present results only for the $R = 8$
 344 case. The first four principal components \mathbf{u}_r are visualized in Fig. 2. In Fig. 1(b) we plot the
 345 reconstruction of the observed clearness index of Fig. 1(a) using just $R = 8$ components. There
 346 certainly is a loss of information when we project the high-dimensional description to the low-
 347 dimensional one. The compression ratio is about 10^5 , however, and we are able to estimate the
 348 error in prediction by the methods described in the previous section.

349 Having constructed the reduced representation of the data, we come to the problem of learn-
 350 ing the evolution of the reduced dynamics described in Sec. 2.4. In Fig. 3 we illustrate the
 351 evolution of $\mathbf{x}(t)$ of Eq. (7). The stars correspond to observations. Notice that since the collection
 352 procedure has both weekly and daily irregularities in its protocol, as described in Sec. 2.1, the
 353 data in Fig. 3 show occasional temporal gaps. As a result, of the 822 observations of $\mathbf{x}(t)$, only
 354 252 remain in \mathcal{D}_R of Eq. (24) and can, therefore, be used to train the recursive GP of Eq. (22).
 355 The reason is that we can only use triplets of three consecutive measurements with a 30-minute
 356 lag, and the gaps in data imply that not all observations are part of such triplets. Such a limited
 357 number of observations in \mathcal{D}_R constrains the number of degrees of freedom in the statistical

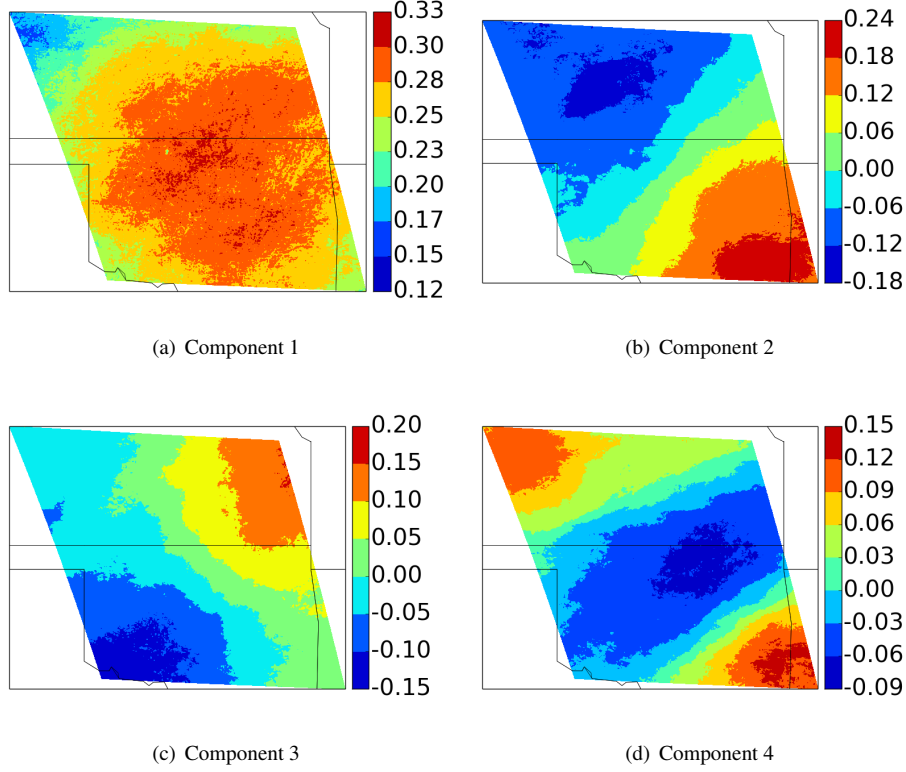


Figure 2: Principal components ($\text{vec}^{-1}(\mathbf{u}_r)$) 1 to 4.

358 model and consequently limits the size of the reduced dimension R .

359 The covariance function we are using in Eq. (27) has the form

$$k(\mathbf{z}, \mathbf{z}'; \boldsymbol{\theta}_r) = v_{r,0} + \sum_{i=1}^{2R} v_{r,i} z_i z'_i + k_0(\mathbf{z}, \mathbf{z}'; \boldsymbol{\theta}_{r,0}) + \sigma_r^2 \delta(\mathbf{z} - \mathbf{z}'), \quad (41)$$

360 where z_i and z'_i are the i -th component of \mathbf{z} and \mathbf{z}' , respectively, $\boldsymbol{\theta}_r = (v_{r,0}, \dots, v_{r,2R}, \boldsymbol{\theta}_{r,0})$, $\sigma_r, v_{r,i} \geq$
361 0 , and k_0 is a classic covariance function (see [34, Ch. 4]). Here, $r = 1, 2, \dots, R$ denotes the index
362 of the reduced components whose evolution is modeled. The first term in Eq. (41) captures a
363 constant bias. The second, nonstationary, term captures a structured linear dependence between
364 \mathbf{z} and \mathbf{x} . The third term captures any nonlinear dependence. The second term can be interpreted
365 as the effect of a mean 0 prior on a linear mean term [34, Ch. 2.7]. The last term models the
366 possible effects of measurement noise (though we anticipate its effect to be small, since the vector
367 \mathbf{x} is spatially filtered by the projection on the principal component). The covariance function
368 forms $k_0(\cdot, \cdot; \cdot)$ we have experimented with include the squared exponential, the exponential and
369 several covariances belonging to the matern class. None of these choices yielded a significant
370 improvement in the final forecasting capabilities of the model. Therefore, we present only those

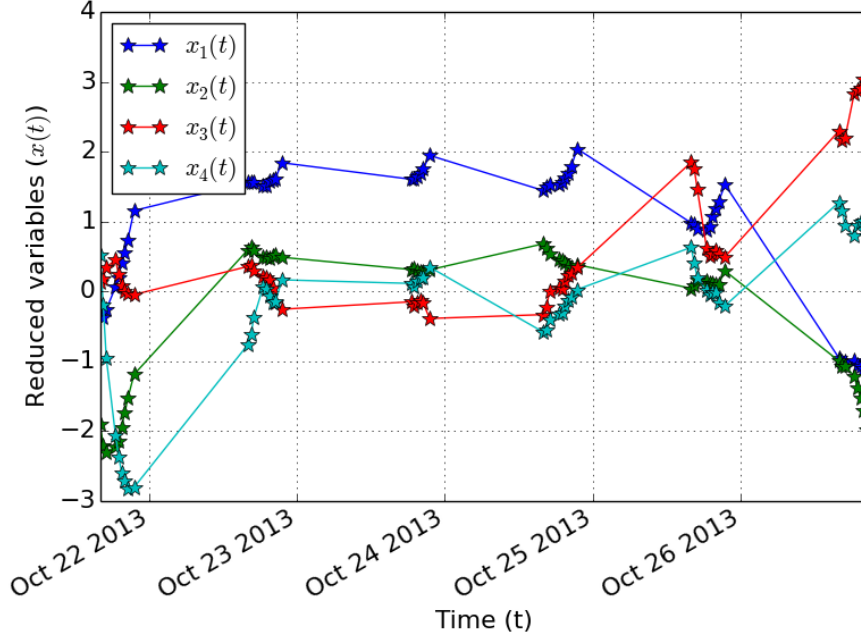


Figure 3: The first four reduced variables $x_1(t), \dots, x_4(t)$. The time is in UTC, and the stars indicate observations.

371 results obtained by the exponential covariance function whose form is

$$k_0(\mathbf{z}, \mathbf{z}', \boldsymbol{\theta}_{r,0}) = s_r^2 \exp \left\{ - \sum_{i=1}^{2R} \frac{|z_i - z'_i|}{\ell_{r,i}} \right\}, \quad (42)$$

372 where $\boldsymbol{\theta}_{r,0} = (s_r, \ell_{r,1}, \dots, \ell_{r,2R})$. The parameters s_r and $\ell_{r,i}$ may be interpreted as the nonlinear
 373 part of the signal strength and its length scale with respect to the input variable z_i , respectively.
 374 The number of parameters per r is $|\boldsymbol{\theta}_r| = 4R + 3$, and the total number of parameters is $R(4R + 3)$.
 375 For $R = 8$ this translates to 35 parameters per r that need to be inferred from $\mathcal{D}_{R,r}$ of Eq. (26). It
 376 now becomes apparent that the low number of observations (about 250) we have at our disposal
 377 does not allow for choosing a large R .

Table 1: The inferred parameters of the dynamic model. The ‘-’ stand for values that are less than 10^{-2} , i.e., essentially zero. The ‘*’ stand for length scales that were estimated to be larger than 10 and, therefore, unimportant.

Parameter	1	2	3	4	5	6	7	8
$v_{r,0}$	-	-	-	-	-	-	-	-
$v_{r,1}$	0.12	-	-	-	-	-	-	-
$v_{r,2}$	-	0.24	-	-	-	-	-	-
$v_{r,3}$	-	-	0.35	-	0.01	-	-	-
$v_{r,4}$	-	-	-	0.20	-	-	-	-
$v_{r,5}$	-	-	-	-	0.32	-	-	-

Continued on next page

Table 1 – Continued from previous page

Parameter	1	2	3	4	5	6	7	8
$v_{r,6}$	-	-	-	-	-	0.24	-	-
$v_{r,7}$	-	-	-	-	-	-	0.35	-
$v_{r,8}$	-	-	-	-	-	-	-	0.30
$v_{r,9}$	1.72	-	-	-	-	-	-	-
$v_{r,10}$	-	0.89	-	-	-	-	-	-
$v_{r,11}$	-	-	2.39	-	0.01	-	-	-
$v_{r,12}$	-	-	-	2.01	-	-	-	-
$v_{r,13}$	-	-	-	-	2.25	-	-	-
$v_{r,14}$	-	-	-	-	-	2.02	-	-
$v_{r,15}$	-	-	-	-	-	-	4.27	-
$v_{r,16}$	-	-	-	-	-	-	-	2.08
s_r^2	0.01	0.01	0.02	0.02	0.04	0.04	0.07	0.06
$\ell_{r,1}$	0.17	*	*	*	*	*	*	1.20
$\ell_{r,2}$	2.24	2.19	*	*	1.22	0.44	*	1.40
$\ell_{r,3}$	*	*	*	*	0.87	*	*	2.02
$\ell_{r,4}$	*	*	2.69	*	5.99	*	*	3.91
$\ell_{r,5}$	*	*	*	3.79	*	1.37	*	*
$\ell_{r,6}$	*	*	*	0.21	*	*	*	*
$\ell_{r,7}$	*	*	*	*	*	*	1.68	*
$\ell_{r,8}$	*	*	*	3.08	0.96	*	0.98	0.46
$\ell_{r,9}$	0.11	0.26	1.16	1.26	*	0.78	*	*
$\ell_{r,10}$	*	1.09	*	*	*	0.89	0.91	*
$\ell_{r,11}$	*	1.65	0.31	3.24	0.80	*	*	*
$\ell_{r,12}$	1.52	0.63	*	*	*	*	2.40	0.96
$\ell_{r,13}$	*	*	*	*	0.29	*	8.32	*
$\ell_{r,14}$	*	*	5.70	0.39	*	0.32	4.07	4.16
$\ell_{r,15}$	*	*	1.50	4.14	*	*	0.75	*
$\ell_{r,16}$	*	*	0.59	6.07	*	*	*	0.37
σ_r^2	-	-	-	-	-	-	-	-

378 The optimization problem of each r (see Eq. (31)) is solved by using the BFGS algorithm
 379 [36]. To accommodate the nonnegative constraints, we use an exponential change of variable.
 380 The results are shown in Table 1. The ‘-’ symbols for the $v_{r,i}$ and the σ_r^2 correspond to values that
 381 are smaller than 10^{-2} and, therefore, essentially zero. Furthermore, the ‘*’ symbols for the $\ell_{r,i}$
 382 scale parameters indicate that their corresponding values were greater than 10 and, hence, tend
 383 to affect the model significantly less than those with smaller values. Both these choices make
 384 dependence patterns more visible. We notice that the bias terms $v_{r,0}$ are effectively zero. This
 385 result was expected since $\mathbf{x}(t)$ are the projections of quantities with mean removed (21). We
 386 also notice that given the size of the parameters s_r^2 , the most important effect observed originates
 387 from the nonstationary component of our covariance function of Eq. (41). This is the part that
 388 captures the linear part of the correlations between $\mathbf{z}(t)$ and $\mathbf{x}(t)$. However, the nonlinear part,
 389 coming from k_0 of Eq. (42), is also considerable and accounts for about 3% to 20% of total
 390 signal strength as measured by $s_r / \sqrt{\sum_{i=0}^{2R} v_{r,i} + s_r^2}$ (though that is not a ratio of the variances
 391 of the components, which has a far more complicated expression). This effect becomes more

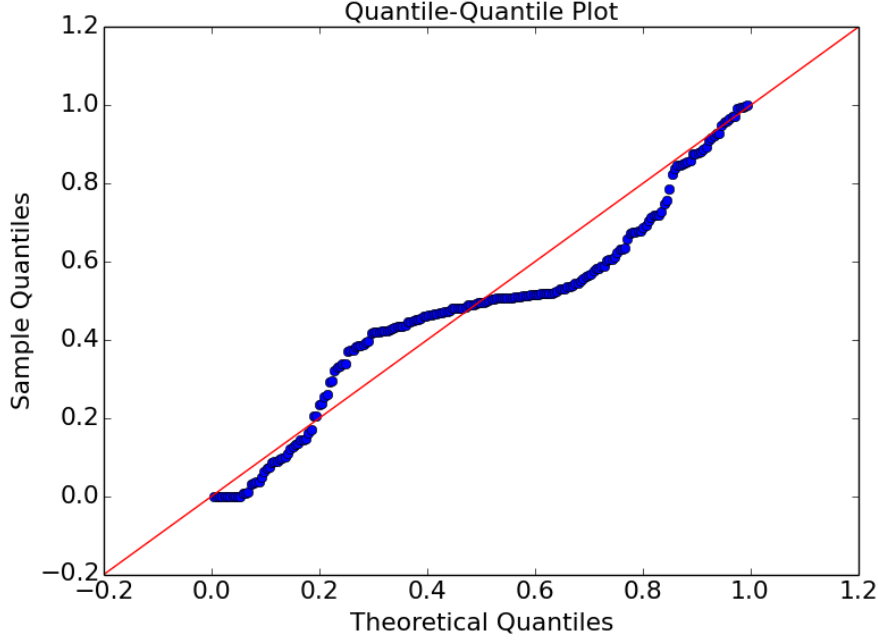


Figure 4: Quantile plot for one-step-ahead probabilistic predictions.

392 important as we go to higher FA coefficients. Another important point to notice is that there is
 393 no significant exchange of energy between the linear parts of the signal. That is, $x_r(t)$ depends
 394 linearly only on $z_r(t) = x_r(t-1)$ and $z_{r+R}(t) = x_r(t-2)$. As expected, the dependence on $x_r(t-1)$
 395 is larger than on $x_r(t-2)$. An explanation of this phenomenon appears in Fig. 6. We see that the
 396 scatterplots of $x_i(t)$ and $x_j(t-1)$ for various values of $j \neq i$ appear to be null plots: essentially
 397 indistinguishable from the plot of independent random variables. This phenomenon is essentially
 398 the same for other $i \neq j$ cases not displayed and this explains the insignificant $v_{r,i}$ coefficients
 399 for $r \neq i$. The length scale parameters $\ell_{r,i}$ of the nonlinear part are interpreted in a reciprocal
 400 manner to $v_{r,i}$. In particular, the larger the length scale, the less important the interaction. As for
 401 the linear part, again the main effect is that $x_r(t)$ depends mostly on $x_r(t-1)$ and $x_r(t-2)$, though
 402 here there is a significant number of cross-component dependence coefficients that are not zero.
 403 The parameter σ_r^2 was found to be close to zero in all cases. From the first plot in Fig. 6, which
 404 again is consistent with $x_i(t)$ by $x_i(t-1)$ plots other than the displayed $i = 1$, we see that the noise
 405 levels are low and are likely to be masked by the other variance components. The small value
 406 of σ_r does not mean that the uncertainty in our predictions will be low (see Eq. (36) for the way
 407 that σ_r^2 affects the predictive variance). Even if these parameters were perfectly known and σ_r
 408 were 0, the uncertainty could go to 0 only in the limit of infinite data.

409 We can now make forecasts for an arbitrary number of steps ahead. The predictive distri-
 410 bution of our model is non-Gaussian because the reduced dynamics are non-linear. However,
 411 it is still possible to test its ability to capture the observed variation of the solar irradiance by
 412 employing numerical methods. In lack of a better alternative, we focus our attention to the solar
 413 irradiance at the center of the computational patch. Consider the clearness indices $C(t-1)$ and

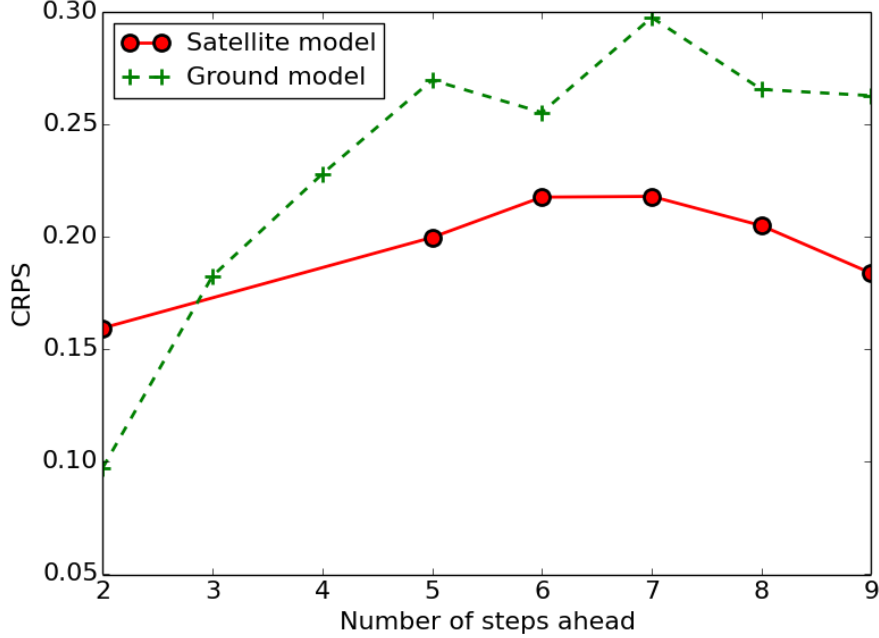


Figure 5: Comparison of the averaged continuous ranked probability score (CRPS), as a function of the number of steps ahead of the forecast, for the satellite (red) and the ground (green) model.

414 $\mathbf{C}(t - 2)$ we observe at times $t - 1$ and $t - 2$, respectively. Let $I_{\text{center, model}}(t)$ and $I_{\text{center, obs}}(t)$ the
 415 predicted and the observed solar irradiance, respectively, at the center of the computational patch
 416 at time t . Consider the cumulative distribution function (CDF) of $I_{\text{center, model}}(t)$:

$$F_t(I) = P[I_{\text{center, model}}(t) \leq I]. \quad (43)$$

417 If our model predicted the observations perfectly, that is if $I_{\text{center, obs}}(t)$ was indeed generated from
 418 our model, then $F_t(I_{\text{center, obs}}(t))$ would be distributed uniformly between 0 and 1 independently of
 419 t . Unfortunately, the CDF $F_t(I)$ is not known analytically, so we need to approximate it. This is
 420 accomplished by obtaining $S = 1028$ samples of $I_{\text{center, model}}(t)$, $\{I_{\text{center}}^{(s)}(t)\}_{s=1}^S$, from the predictive
 421 distribution conditional on $\mathbf{C}(t - 1)$ and $\mathbf{C}(t - 2)$ following the instructions of Sec. 2.5. Using
 422 these samples, we approximate $F_t(I)$ via the empirical CDF,

$$\hat{F}_{t,S}(I) = \frac{1}{S} \sum_{s=1}^S 1_{\{I_{\text{center, model}}^{(s)}(t) < I\}}(I), \quad (44)$$

423 where $1_A(x)$ is the indicator function of the set A , i.e. $1_A(x) = 1$ if $x \in A$ and 0 otherwise.
 424 Now, we can compare the quantiles of $\hat{F}_{t,S}(I_{\text{center, obs}}(t))$ with the quantiles of the uniform distri-
 425 bution between 0 and 1. We do this for all observations we have in our validation set, i.e., for
 426 all useful satellite measurements of the solar irradiation between February 26, 2014, and May
 427 5, 2014. The result is shown in Fig. 4. The probability distributions agree in the sense that the
 428 measures induced by them are compatible: intervals of certain length on the x axis (which are

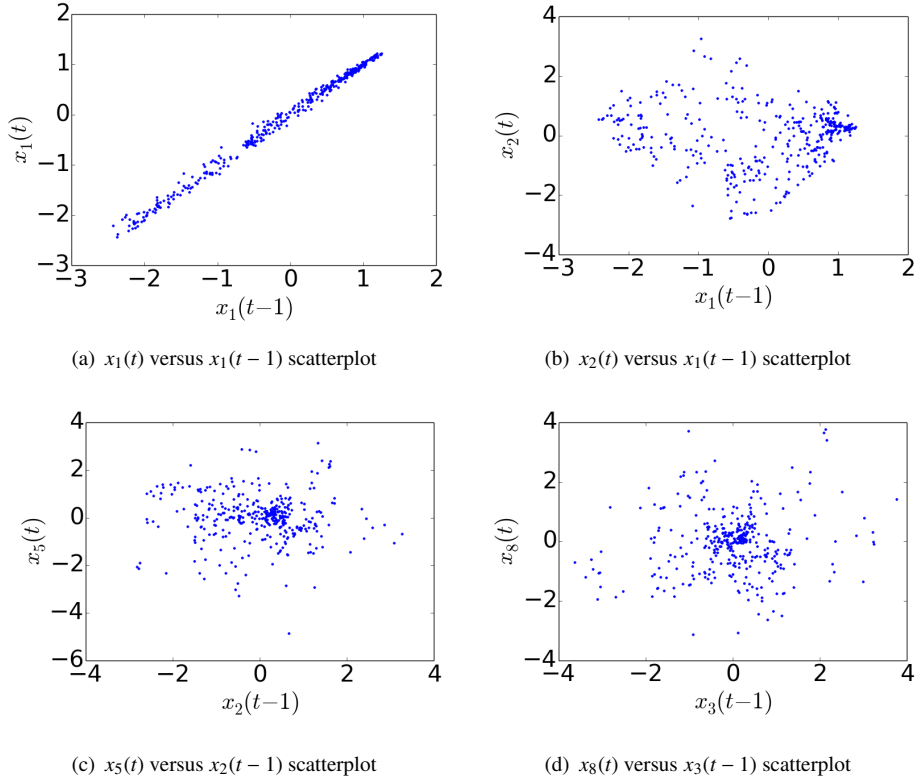


Figure 6: Scatterplots of $x_i(t)$ by $x_j(t-1)$, for $(i, j) = (1, 1), (1, 2), (2, 5), (3, 8)$

429 the nominal lengths of the corresponding prediction intervals) are mapped into comparable inter-
 430 vals. There seems to be a moderate to small amount of misspecification – misfit in the quantile
 431 plot in Fig. 4 – that does not disappear as the amount of data would increase, a situation not
 432 uncommon in atmospheric and environmental models. The previous observation indicates that,
 433 however, the prediction intervals, while not having the prescribed nominal coverage would have
 434 one comparable to it, which means that the variability would be reasonably well captured.

435 At this point, we would like to compare the forecast of our satellite model to that of a model
 436 that is built using only pointwise ground observations. In particular, we would like to know if
 437 making use of the satellite observations does improve the forecast quality and to what extent.
 438 Toward this end, we construct a solar irradiation model using a recursive GP model trained
 439 only on ground observations. This model is equivalent to our satellite model when using only
 440 one component which is obtained directly from the data (i.e., no-dimensionality reduction is
 441 necessary). We will refer to it as the *ground model*. It has a total of 7 parameters which we
 442 learn using 350 ground observations. Next, we compute the continuous ranked probability score
 443 (CRPS) [37] of both the satellite and ground models and compare them to each other. The
 444 CRPS can be thought of as an extension of the mean absolute error (MAE) score to probabilistic
 445 forecasts. It is essentially a measure of the closeness of the predictive distribution to the one
 446 generated by the data. The smaller it is, the more accurately the model captures the underlying

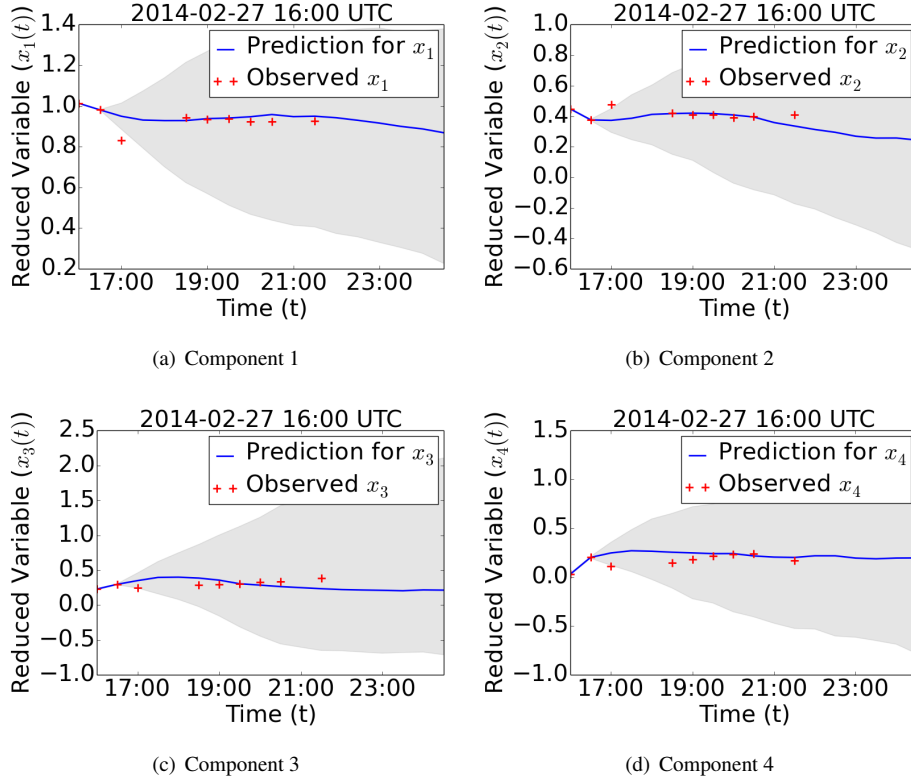


Figure 7: February 27, 2014 (sunny day): Median and 95% prediction interval for the reduced variables $\mathbf{x}(t)$ on February 27, 2014.

447 stochastic process. Extending a little bit the notation of the previous paragraph, we take $S = 1028$
 448 samples of $I_{\text{center}}(t, t')$, $\{I_{\text{center},a}^{(s)}(t, t')\}$, $s = 1, \dots, S$ —where a stands for either the “satellite” or
 449 “ground” model— from the predictive distribution conditional on $\mathbf{C}(t')$ and $\mathbf{C}(t' + 1)$, if dealing
 450 with the satellite model, or $I_{\text{center,ground,obs}}(t')$ and $I_{\text{center,ground,obs}}(t' + 1)$, if dealing with the ground
 451 model, following the instructions of Sec. 2.5. The discrete version of CRPS is given by [37]:

$$\begin{aligned} \text{CRPS}_{a,S}(t, t') &= \frac{1}{S} \sum_{s=1}^S |I_{\text{center},a}(t, t')^{(s)} - I_{\text{center},a,obs}(t)| \\ &\quad - \frac{1}{2S^2} \sum_{s=1}^S \sum_{q=1}^S |I_{\text{center},a}(t, t')^{(s)} - I_{\text{center},a}(t, t')^{(q)}|. \end{aligned}$$

452 Let n be the number of steps ahead for which we wish to make a prediction of the solar irradiation.
 453 We will just refer to n as the “number of steps ahead”. In Fig. 5, we plot the *mean* of the CRPS
 454 —the mean is over the part of the validation set recorded at 16:00UTC— as a function of the
 455 number of steps ahead, n . That is, we plot the mean $\text{CRPS}_{a,S}(n)$ of $\text{CRPS}_{a,S}(t' + n, t')$ for t'
 456 in the validation set corresponding to time 16:00UTC. We observe that the ground model does
 457 better at the very beginning, but that its performance gradually deteriorates. This is due to the

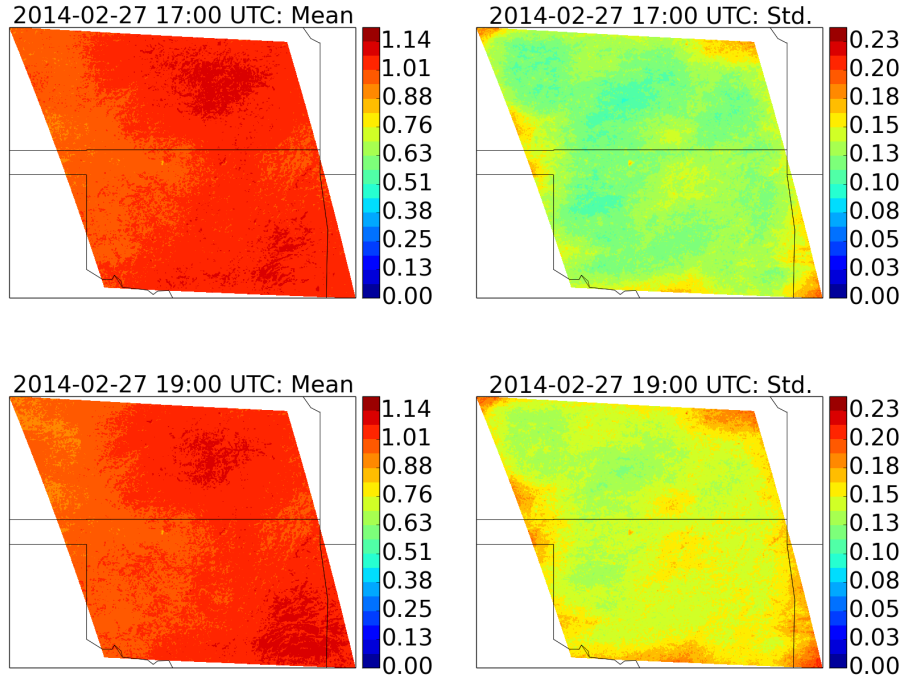


Figure 8: February 27, 2014 (sunny day): Starting with two observations on February 27, 2014, at 16:00 and 16:30 UTC, we show the forecasts of the clearness index for 17:00 and 19:00 UTC. The left column shows the mean forecast and the right column the standard deviation.

458 fact that the satellite model includes additional errors rising from the dimensionality reduction
 459 step. For longer forecasts, however, the satellite model does significantly better. The remarkable
 460 fact is that its CRPS performance remains more or less constant over time. These conclusions
 461 are reached again via Figs. 11 and 16 and the corresponding discussion in the text.

462 Now, we show results for two representative days: February 26, 2014, which was a sunny
 463 day, and March 1, 2014, which was a cloudy day. To initiate our forecast, we observe the clear-
 464 ness index at 16:00 and 16:30 UTC. Then, we make predictions for 8 hours ahead by sampling
 465 Eq. (36) recursively as explained in Alg. 1. For the computation of the quantiles we present here,
 466 we use 1,000 such samples.

467 *Forecast for February 27, 2014 (sunny day).* Figure 7 depicts the evolution of the median and
 468 95% prediction intervals of our prediction for the first four components of $\mathbf{x}(t)$. We notice that
 469 the observed values always fall within the error bars. Fig. 8 shows the evolution of the mean and
 470 the standard deviation of $C(t)$, respectively. Since the day is sunny, the mean remains relatively
 471 constant. Notice that the standard deviation of $C(t)$ does increase with time but not very fast. In
 472 Fig. 9 we compare the observed $C(t)$ with our mean predictions at three different times. Notice
 473 that high-frequency features are not captured properly. This result is expected, however, since
 474 the model is built on the projected $C(t)$. In Fig. 10, we compare the error between projection
 475 of $C(t)$ and our mean predictions, with the forecasted standard deviation. We notice that the

476 error is indeed well bounded by the standard deviation, which shows that our model is at least
477 conservative. In Fig. 11 we compare our point predictions for the solar irradiation at Lamont, OK,
478 with the predictions of the ground model. We see that the mean prediction is considerably better
479 for our model compared with the one that uses only data from the prediction site, particularly
480 a few hours after the start of the forecast window. Moreover, the standard deviation of the
481 global data approach is smaller. We note here that a discrepancy exists between the satellite
482 and the ground observations. This is expected, since the satellite solar irradiation is averaged
483 over a square kilometer while the ground measurement is almost pointwise. Overall we find that
484 the uncertainty margins are improved and that the predictions are vastly improved when using
485 satellite data compared with only site data.

486 *Forecast for March 1, 2014 (cloudy day).* Figure 12 depicts the evolution of the median and
487 95% prediction intervals of our prediction for $\mathbf{x}(t)$. We notice that $\mathbf{x}(t)$ takes more extreme values
488 compared with Fig. 7. As a result, some of the observations and, in particular, those of $x_2(t)$
489 do not always fall within the error bars, though their errors are certainly on the order of the
490 standard deviations. Fig. 13 shows the evolution of the mean and the standard deviation of $\mathbf{C}(t)$,
491 respectively. Here we notice a much larger variation of the mean prediction compared with
492 Figure 8. In addition, the standard deviation increases much more rapidly. In Fig. 14 we compare
493 the observed $\mathbf{C}(t)$ with our mean predictions at three different times. Here, the error is much
494 larger, and the standard deviation captures the order of the larger errors but not their pattern,
495 though it does so for the large feature persistent error. A look at the projected version of the
496 error (see Fig. 15) validates this observation. Another interesting property is that the standard
497 deviation increases faster near the boundaries of the patch than at the center. The cause is that
498 effects near the boundary depend also on the clearness index field that is outside our small patch.
499 In Fig. 16 we compare our point predictions for the solar irradiation at Lamont, OK, with that of
500 a recursive GP model based only on ground observations. The mean prediction has smaller error
501 for the satellite data compared with the site data, whereas the standard deviation is significantly
502 smaller after a couple of hours. The performance of our model is reasonable and certainly better
503 than the model based on ground observations. However, sharp variations such as the ramps
504 occurring at 20:00 UTC are not captured. We believe that in order to capture ramp events one
505 has to model the finer frequencies, which is not done in the version of the model. We note here
506 that ramps could not be captured even with $R = 8$, which was the largest model that our limited
507 observations allowed us to train. The discrepancy between the measurements at 16:00 and 16:30
508 UTC and the starting points of our forecasts is due to the fact that with $R = 8$ the clearness index
509 cannot be reconstructed perfectly from the reduced space. Such a discrepancy is present also at
510 16:00 UTC in Fig. 11 but much less pronounced. However, we also note that our pPCA approach
511 reasonably captures that uncertainty.

512 In any case, we again observe a significant improvement in forecast quality in comparison to
513 the model that uses only site data, though in this case the improvement is more pronounced in
514 the variance estimation.

515 4. Conclusions

516 We presented a probabilistic forecast model for solar irradiation based on satellite observa-
517 tions. The approach is based on separate modeling of the clear sky model and of the clearness
518 index, the ratio of irradiation and the clear sky model. Since good clear sky models exist, only
519 the clearness index needs to be modeled statistically. The resulting model yields probabilistic

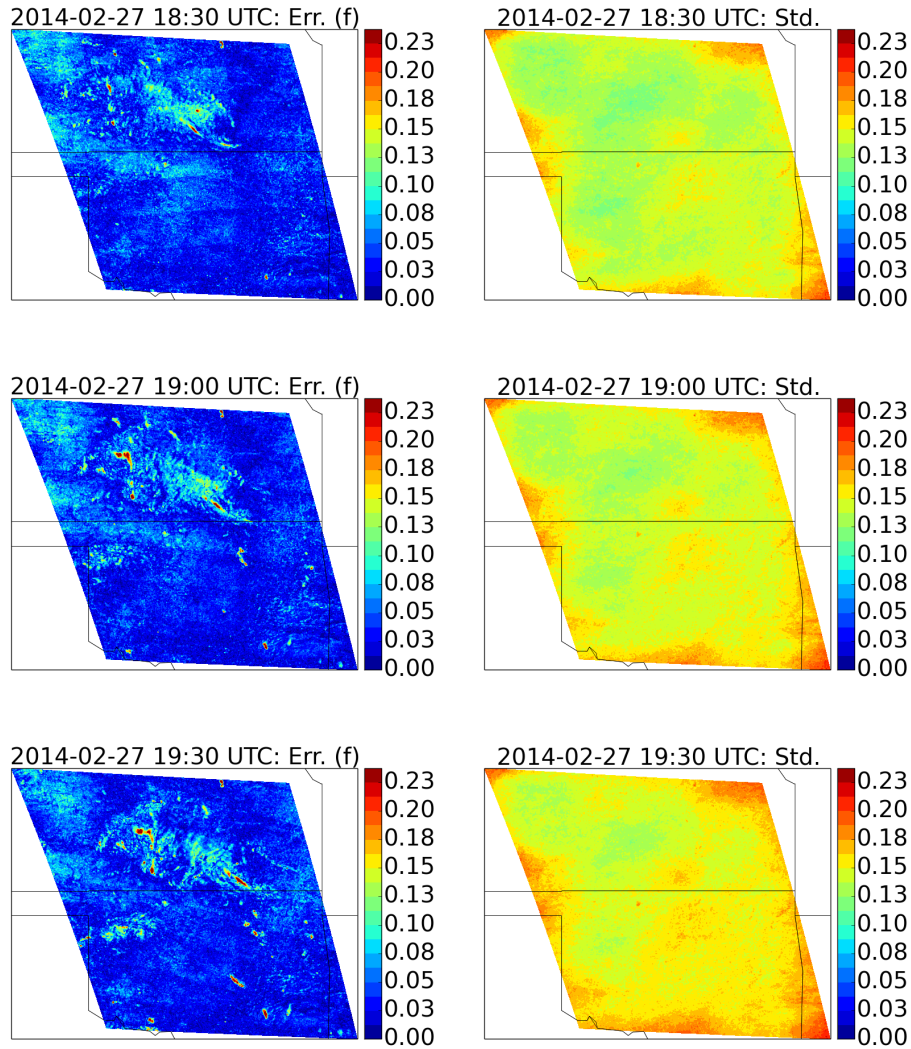


Figure 9: February 27, 2014 (sunny day): Comparison of the real absolute error in the prediction (left) with the standard deviation of our forecast (right) on February 27, 2014.

520 estimates of the irradiation field every thirty minutes starting from consecutive satellite measure-
 521 ments. The model is based on a probabilistic PCA approach to reduce the forecast dimension,
 522 followed by a Gaussian process approach to learn its dynamics. Using ground data from Lamont,
 523 OK, we demonstrate that the model results in better predictions and uncertainty estimates for the
 524 solar radiation compared with models that use only ground data at the prediction site, which is
 525 the prevailing forecast method for solar energy operations. Moreover, since the data are avail-
 526 able over the entire United States, forecasts can be produced at any site in the United States with
 527 (what we anticipate) to be comparable accuracy, without needing ground measurements. This

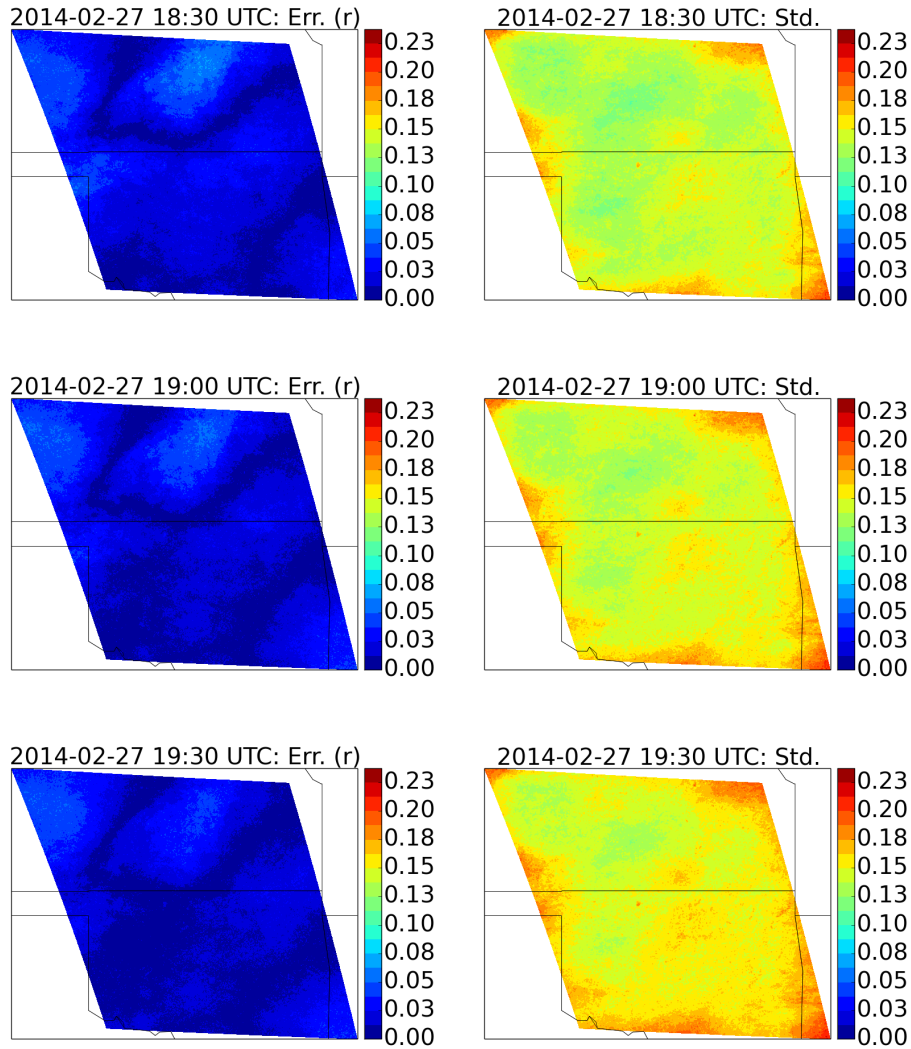


Figure 10: February 27, 2014 (sunny day): Comparison of the absolute projected error in the prediction (left) with the standard deviation of our forecast (right) on February 27, 2014.

528 may prove important for very distributed solar energy systems.

529 This initial investigation can be improved in multiple ways, not the least of which is access
 530 to more data as the AVHRR instrument acquires more measurements. This extra data will allow
 531 us to consider more components and richer covariance functions, and, in particular, to model the
 532 larger space frequencies which we currently do not do as well as the satellite resolution errors.
 533 As a result, we may be able to represent higher-fidelity spatial features such as the ramps that we
 534 cannot capture well for the March 1, 2014, forecast.

535 We note, however, that in some applications solar energy sources are spread over large (and

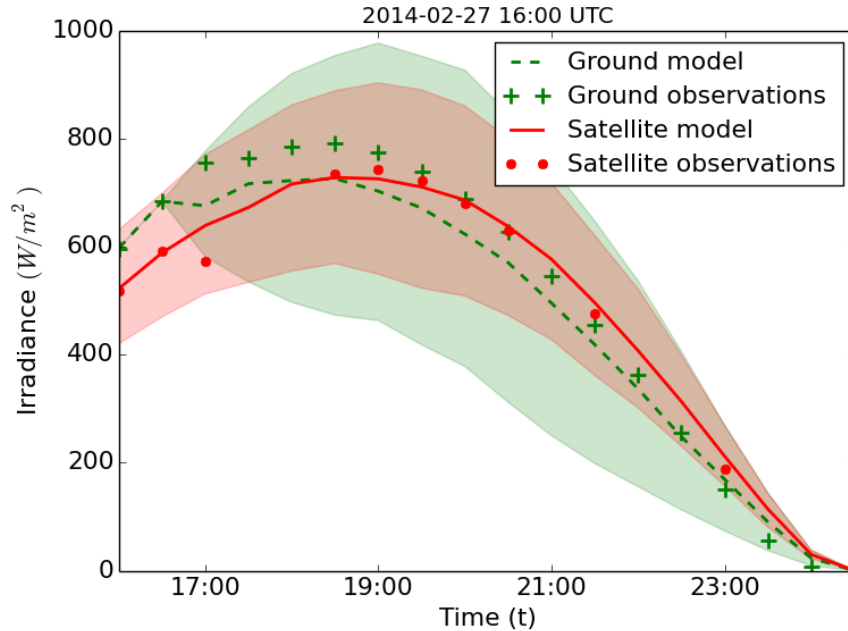


Figure 11: 27 February, 2014 (sunny day): Comparison of our point forecast at Lamont, OK, with a recursive Gaussian process model based solely on ground observations.

perhaps well-chosen) areas. In such cases, the fact that we capture well the low-frequency components of the error (interpreted here as the projection on the space of the principal components, as is seen in Figs. 10 and 15 indicates that spatially aggregated statistics of available solar power are likely to be well captured. However, we would have no way of obtaining that much ground data to validate the predictions. Nevertheless, looking at the plots of the ground irradiation and satellite radiation data, we find a good agreement, which indicates that this is likely to be true. Moreover, one would need to choose a pattern for the solar plants, which requires obtaining some realistic configurations; this will be the subject of future analyses.

Another interesting direction is to combine these measurements with weather forecast information, which –by itself– has difficulties in representing accurately irradiation information because of poor high-resolution cloud forecasting. However, since weather forecast may capture certain large-scale trends (e.g. wind), such information when combined with this approach conceivably may result in better forecasts.

Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-06CH11357. We gratefully acknowledge the use of the Blues cluster in the Laboratory Computing Resource Center at Argonne National Laboratory. We thank Christine Molling from the Cooperative Institute for Meteorological Satellite Studies, Space Science and Engineering Center, at the University of Wisconsin-Madison for providing access to CLAVRx

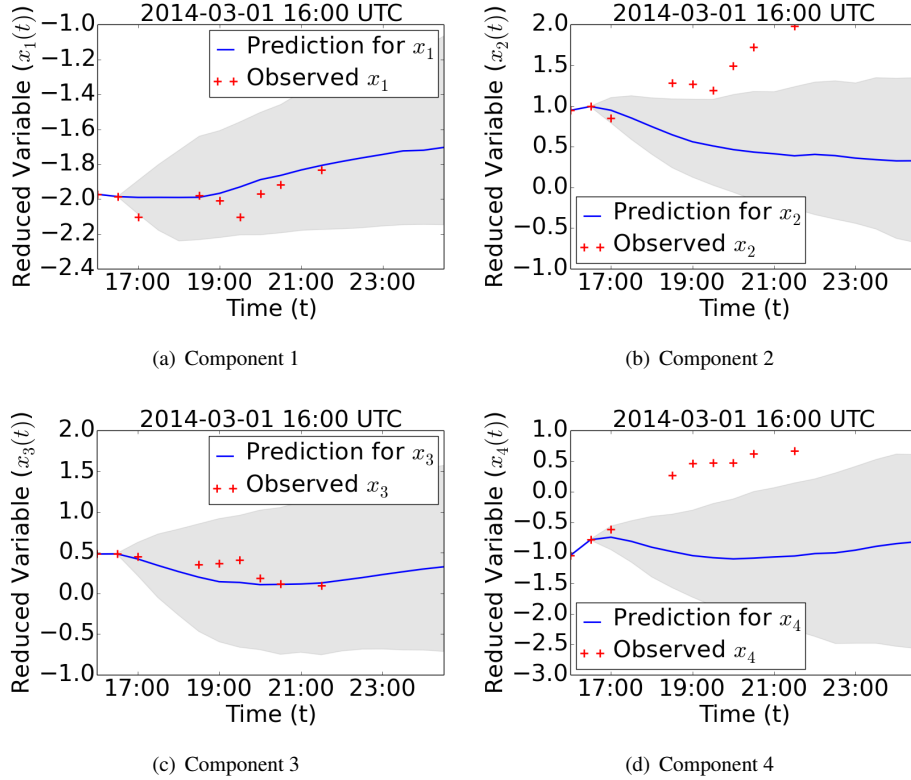


Figure 12: March 1, 2014 (cloudy day): Median and 95% prediction interval for the reduced variables $\mathbf{x}(t)$ on March 1, 2014.

555 Satellite data, and Edwin Campos, Research Meteorologist, Argonne National Laboratory for
 556 his help in interpreting them. We also acknowledge the Atmospheric Radiation Measurement
 557 (ARM) Program for providing access to the SGP ground irradiance measurements.

558 Appendix A. Details of the clear sky model

559 As discussed in Sec. 2.2, we use Ineichen's model [12] as our clear sky model. Here we
 560 briefly discuss the details of its implementation that pertain to this work. For the actual model,
 561 the reader may consult the original reference. Ineichen's model requires the solar zenith, the
 562 extraterrestrial irradiation, the surface elevation, the water vapor content, and the aerosol optical
 563 depth at 700 nm.

564 The solar zenith is computed by using the Solar Position Algorithm (SPA) [38]. The extrater-
 565 restrial irradiation is modeled as [39]

$$I_{\text{ext}}(t) = I_0 \left(\frac{R(t)}{R_{\text{av}}} \right)^2, \quad (\text{A.1})$$

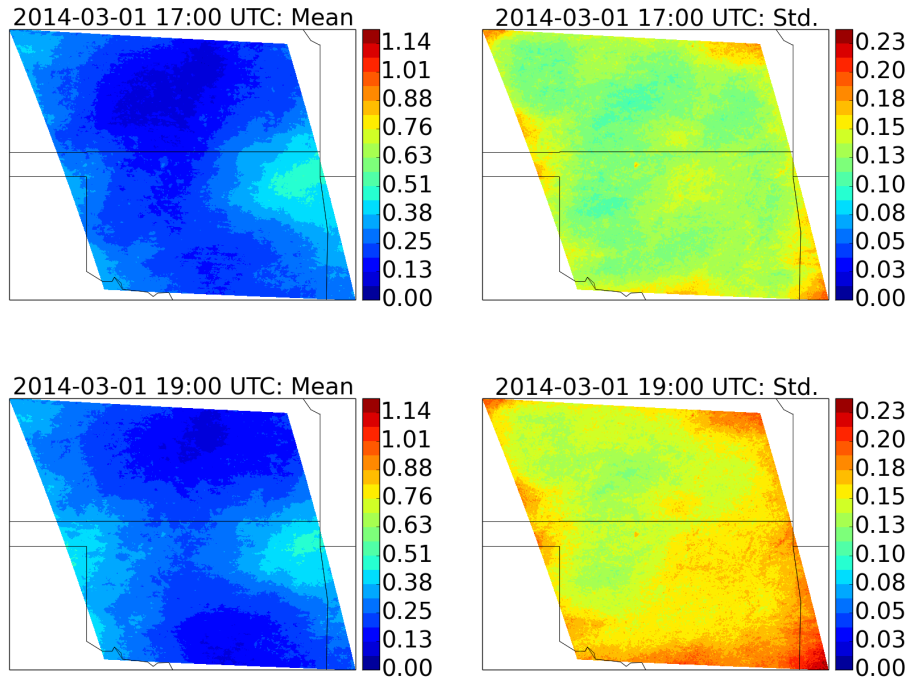


Figure 13: March 1, 2014 (cloudy day): Starting with two observations on March 1, 2014, at 16:00 and 16:30 UTC, we show the forecasts of the clearness index for 17:00 and 19:00 UTC. The left column shows the mean forecast and the right column the standard deviation.

566 where $R(t)$ is the Earth-Sun distance at time t , $I_0 = 1361 \text{ W/m}^2$ and $R_{\text{av}} = 1.0000010178 \text{ AU}$ is
 567 the average Earth-Sun distance. The Earth-Sun distance is also computed by using SPA. Since
 568 the sensitivity of the model to the surface elevation is small and the patch of the CONUS scan
 569 we are concentrating on is relatively flat, we take the surface elevation to be constant and equal
 570 to 320 m. The water vapor content and the aerosol optical depth at 700 nm are also taken to be
 571 constant at 0.3 cm and 0.15, respectively.

572 References

- 573 [1] G. Box, G. Jenkins, G. Reinsel, Time Series Analysis: Forecasting and Control, Wiley Series in Probability and
 574 Statistics, Wiley, 2008.
 575 URL <http://books.google.com/books?id=1JnnPQAACAAJ>
- 576 [2] J. Boland, Time series modeling of solar radiation, in: V. Badescu (Ed.), Modeling Solar Radiation at the Earth's
 577 Surface, Springer Berlin Heidelberg, 2008, pp. 283–312. doi:10.1007/978-3-540-77455-6_11.
 578 URL http://dx.doi.org/10.1007/978-3-540-77455-6_11
- 579 [3] J. Kleissl, Solar Energy Forecasting and Resource Assessment, Academic Press, 2013.
- 580 [4] A. Mellit, Artificial intelligence technique for modelling and forecasting of solar radiation data; a review, Int. J.
 581 Artif. Intell. Soft Comput. 1 (1) (2008) 52–76. doi:10.1504/IJAISC.2008.021264.
 582 URL <http://dx.doi.org/10.1504/IJAISC.2008.021264>
- 583 [5] C. W. Chow, B. Urquhart, M. Lave, A. Dominguez, J. Kleissl, J. Shields, B. Washom, Intra-hour forecasting
 584 with a total sky imager at the {UC} san diego solar energy testbed, Solar Energy 85 (11) (2011) 2881–2893.

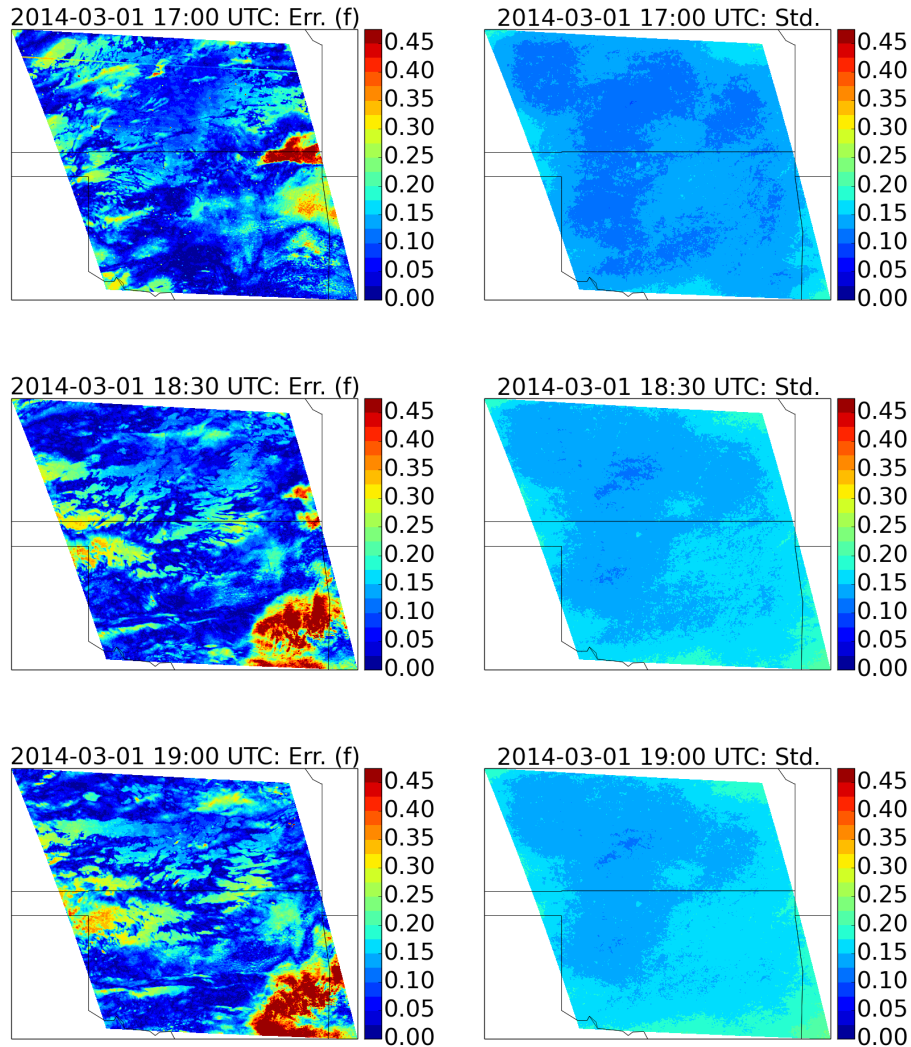


Figure 14: March 1, 2014 (cloudy day): Comparison of real absolute error in the prediction (left) with the standard deviation of our forecast (right) on March 1, 2014.

585 doi:<http://dx.doi.org/10.1016/j.solener.2011.08.025>.
 586 URL <http://www.sciencedirect.com/science/article/pii/S0038092X11002982>
 587 [6] E. Lorenz, D. Heinemann, A. Hammer, Short-term forecasting of solar radiation based on satellite data, in: EuroSun
 588 Conference Freiburg, Germany, 2004, pp. 841–848.
 589 [7] A. Hammer, D. Heinemann, C. Hoyer, R. Kuhlemann, E. Lorenz, R. Miller, H. G. Beyer, Solar energy assessment
 590 using remote sensing technologies, Remote Sensing of Environment 86 (3) (2003) 423–432, urban Remote Sensing.
 591 doi:[http://dx.doi.org/10.1016/S0034-4257\(03\)00083-X](http://dx.doi.org/10.1016/S0034-4257(03)00083-X).
 592 URL <http://www.sciencedirect.com/science/article/pii/S003442570300083X>
 593 [8] J. H. Eto, R. J. Thomas, Computational needs for the next generation electric grid, Department of Energy, (2011).
 594 URL http://energy.gov/sites/prod/files/FINAL_CompNeeds_Proceedings2011.pdf

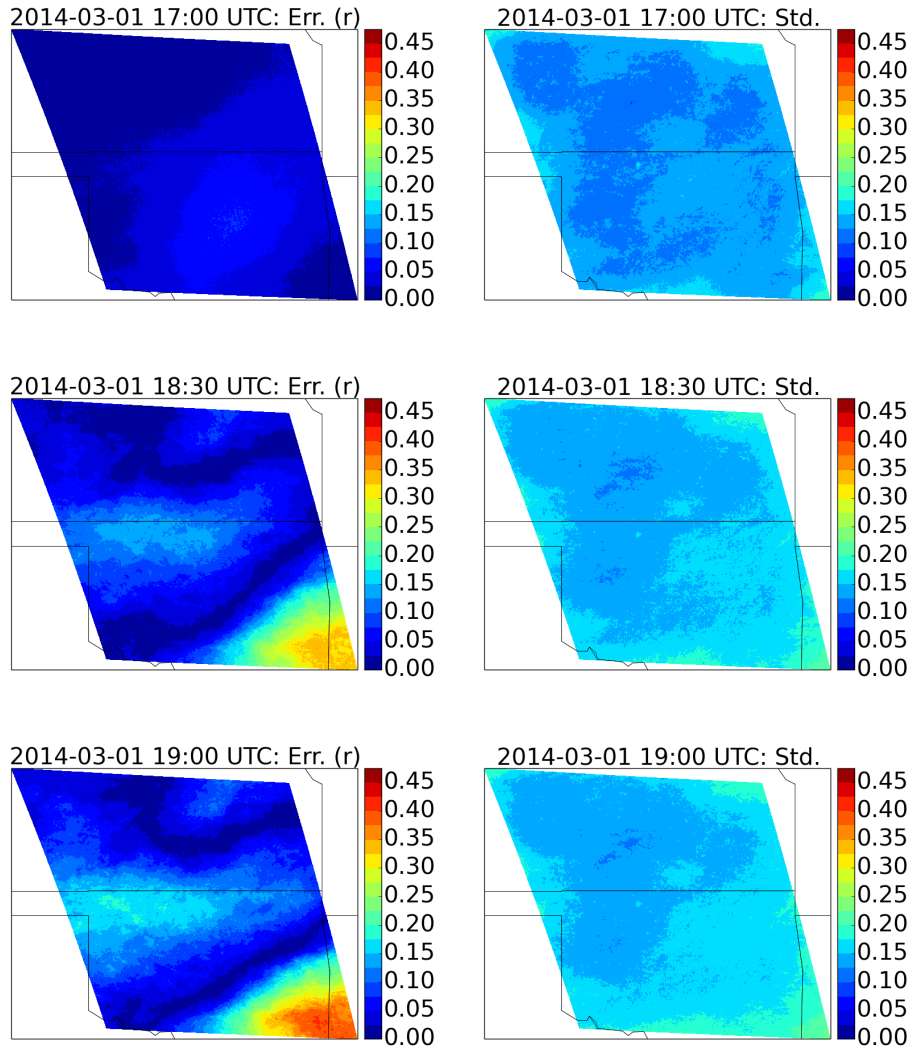


Figure 15: March 1, 2014 (cloudy day): Comparison of absolute projected error in the prediction (left) with the standard deviation of our forecast (right) on March 1, 2014.

- 595 [9] E. M. Constantinescu, V. M. Zavala, M. Rocklin, S. Lee, M. Anitescu, A computational framework for uncertainty
596 quantification and stochastic optimization in unit commitment with wind power generation, *Power Systems, IEEE*
597 *Transactions on* 26 (1) (2011) 431–441.
- 598 [10] I. Laszlo, P. Ciren, H. Liu, S. Kondragunta, J. D. Tarpley, M. D. Goldberg, Remote sensing of aerosol and radiation
599 from geostationary satellites, *Advances in Space Research* 41 (11) (2008) 1882–1893. doi:<http://dx.doi.org/10.1016/j.asr.2007.06.047>.
600 URL <http://www.sciencedirect.com/science/article/pii/S0273117707007156>
- 601 [11] R. E. Bird, R. L. Hulstrom, A simplified clear sky model for direct and diffuse insolation on horizontal surfaces,
602 Tech. Rep. SERI/TR-335-76181, Solar Energy Research Institute (1981).
603
- 604 [12] P. Ineichen, A broadband simplified version of the solis clear sky model, *Solar Energy* 82 (8) (2008) 758–762.




Figure 16: March 1, 2014 (cloudy day): Comparison of our point forecast at Lamont, OK, with a recursive Gaussian process model based solely on ground observations.

605 doi:<http://dx.doi.org/10.1016/j.solener.2008.02.009>.
606 URL <http://www.sciencedirect.com/science/article/pii/S0038092X08000406>
607 [13] R. Mueller, K. Dagestad, P. Ineichen, M. Schroedter-Homscheidt, S. Cros, D. Dumortier, R. Kuhlemann, J. Olseth,
608 G. Piernavieja, C. Reise, L. Wald, D. Heinemann, Rethinking satellite-based solar irradiance modelling: The
609 {SOLIS} clear-sky module, Remote Sensing of Environment 91 (2) (2004) 160 – 174. doi:<http://dx.doi.org/10.1016/j.rse.2004.02.009>.
610 URL <http://www.sciencedirect.com/science/article/pii/S0034425704000690>
611 [14] B. Everitt, An Introduction to Latent Variable Models, Monographs on statistics and applied probability, Chapman
612 and Hall, 1984.
613 URL <http://books.google.com/books?id=2Ug10gAACAAJ>
614 [15] A. Basilevsky, Statistical Factor Analysis and Related Methods: Theory and Applications, Wiley Series in Probabi-
615 lity and Statistics, Wiley, 2009.
616 URL <http://books.google.com/books?id=uvZkWmRTZhUC>
617 [16] D. Bartholomew, M. Knott, I. Moustaki, Latent Variable Models and Factor Analysis: A Unified Approach, Wiley
618 Series in Probability and Statistics, Wiley, 2011.
619 URL <http://books.google.com/books?id=5YxvE1S56DwC>
620 [17] S. Roweis, EM Algorithms for PCA and SPCA, in: in Advances in Neural Information Processing Systems, MIT
621 Press, 1998, pp. 626–632.
622 [18] M. E. Tipping, C. M. Bishop, Probabilistic principal component analysis, Journal of the Royal Statistical Society,
623 Series B 61 (1999) 611–622.
624 [19] K. Pearson, On lines and planes of closest fit to systems of points in space, Philosophical Magazine 2 (1901)
625 559–572.
626 [20] H. Hotelling, Analysis of a complex of statistical variables into principal components, Journal of Educational
627 Psychology 24 (1933) 417–441.
628 [21] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer-Verlag
629 New York, Inc., Secaucus, NJ, USA, 2006.
630

- 631 [22] D. Rubin, D. Thayer, EM algorithms for ML factor analysis, *Psychometrika* 47 (1) (1982) 69–76. doi:10.1007/
632 bf02293851.
633 URL <http://dx.doi.org/10.1007/bf02293851>
- 634 [23] T. P. Minka, Automatic choice of dimensionality for PCA, Tech. Rep. 514 (2000).
- 635 [24] A. Girard, J. Q. Candela, R. Murray-smith, C. E. Rasmussen, Gaussian process priors with uncertain inputs -
636 application to multiple-step ahead time series forecasting, in: *Advances in Neural Information Processing Systems*
637 15, 2002.
- 638 [25] J. M. Wang, D. J. Fleet, A. Hertzmann, Gaussian process dynamical models, in: *Advances in Neural Information*
639 *Processing Systems* 19, MIT Press, 2006, pp. 1441–1448.
- 640 [26] C. K. Liu, A. Hertzmann, Z. Popovic, Learning physics-based motion style with nonlinear inverse optimization,
641 *ACM Trans. Graph* 24 (2005) 1071–1081.
- 642 [27] L. Raskin, E. Rivlin, M. Rudzsky, Using Gaussian process annealing particle filter for 3D human tracking,
643 *EURASIP Journal on Advances in Signal Processing* 2008.
- 644 [28] F. Perez-Cruz, S. Van Vaerenbergh, J. J. Murillo-Fuentes, M. Lazaro-Gredilla, I. Santamaria, Gaussian Processes
645 for Nonlinear Signal Processing: An Overview of Recent Advances, *IEEE Signal Processing Magazine* 30 (4)
646 (2013) 40–50. doi:10.1109/msp.2013.2250352.
647 URL <http://dx.doi.org/10.1109/msp.2013.2250352>
- 648 [29] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learn-
649 ing, MIT Press, Cambridge, MA, USA, 2006.
- 650 [30] S. Conti, A. O’Hagan, U. of Sheffield. Dept. of Probability, Statistics, Bayesian Emulation of Complex Multi-
651 Output and Dynamic Computer Models, Research report (University of Sheffield. Dept. of Probability and Statis-
652 tics), Department of Probability & Statistics, University of Sheffield, 2007.
653 URL <http://books.google.com/books?id=0bQOMwEACAAJ>
- 654 [31] M. Alvarez, N. D. Lawrence, Sparse Convolved Gaussian Processes for Multi-output Regression, in: *NIPS*, 2008,
655 pp. 57–64.
656 URL http://books.nips.cc/papers/files/nips21/NIPS2008_0170.pdf
- 657 [32] I. Bilonis, N. Zabaras, Multi-output local Gaussian process regression: Applications to uncertainty quantification,
658 *Journal of Computational Physics* 231 (17) (2012) 5718–5746. doi:http://dx.doi.org/10.1016/j.jcp.
659 2012.04.047.
660 URL <http://www.sciencedirect.com/science/article/pii/S0021999112002513>
- 661 [33] I. Bilonis, N. Zabaras, B. A. Konomi, G. Lin, Multi-output separable gaussian process: Towards an efficient, fully
662 bayesian paradigm for uncertainty quantification, *Journal of Computational Physics* 241 (0) (2013) 212 – 239.
663 doi:http://dx.doi.org/10.1016/j.jcp.2013.01.011.
664 URL <http://www.sciencedirect.com/science/article/pii/S0021999113000417>
- 665 [34] C. E. Rasmussen, C. K. I. Williams, *Gaussian Processes for Machine Learning* (Adaptive Computation and Machine
666 Learning), The MIT Press, 2005.
- 667 [35] Atmospheric Radiation Measurement (ARM) Climate Research Facility. 1994, updated hourly. Multifilter Ra-
668 diometer (MFR10M). 2013-10-01 to 2014-03-05, 36.605 N 97.485 W: Southern Great Plains (SGP) Cen-
669 tral Facility, Lamont, OK (C1). Compiled by G. Hodges. Atmospheric Radiation Measurement (ARM)
670 Climate Research Facility Data Archive: Oak Ridge, Tennessee, USA. Data set accessed 2014-03-07 at
671 <http://dx.doi.org/10.5439/1025224>.
- 672 [36] M. Avriel, *Nonlinear Programming: Analysis and Methods*, Prentice-Hall, Englewood Cliffs, NJ, 1976.
673 URL <http://www.ams.org/mathscinet-getitem?mr=58:9264>
- 674 [37] E. Gneiting, T. Gneiting, V. Berrocal, N. Johnson, The continuous ranked probability score for circular variables and
675 its application to mesoscale forecast ensemble verification, *Quarterly Journal of the Royal Meteorological Society*
676 132 (621C) (2006) 2925–2942.
- 677 [38] I. Reda, A. Andreas, Solar position algorithm for solar radiation applications, *Solar Energy* 76 (5) (2004) 577 –
678 589. doi:http://dx.doi.org/10.1016/j.solener.2003.12.003.
679 URL <http://www.sciencedirect.com/science/article/pii/S0038092X0300450X>
- 680 [39] G. W. Paltridge, C. M. R. Platt, *Radiative processes in meteorology and climatology* / by G. W. Paltridge and C. M.
681 R. Platt, Elsevier Scientific Pub. Co, Amsterdam, 1976.

Government License The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.