Gaussian Conditional Random Fields for Aggregation of Operational Aerosol Retrievals

Nemanja Djuric, Vladan Radosavljevic, Zoran Obradovic, and Slobodan Vucetic

***Abstract*—We present a Gaussian conditional random field model for the aggregation of aerosol optical depth (AOD) retrievals from multiple satellite instruments into a joint retrieval. The model provides aggregated retrievals with higher accuracy and coverage than any of the individual instruments while also providing an estimation of retrieval uncertainty. The proposed model finds an optimal temporally smoothed combination of individual retrievals that minimizes the root-mean-squared error of AOD retrieval. We evaluated the model on five years (2006–2010) of satellite data over North America from five instruments (Aqua and Terra MODIS, MISR, SeaWiFS, and the Ozone Monitoring Instrument), collocated with ground-based Aerosol Robotic Network ground-truth AOD readings, clearly showing that the aggregation of different sources leads to improvements in the accuracy and coverage of AOD retrievals.**

***Index Terms*—Aerosol optical depth (AOD), data aggregation, Gaussian conditional random fields (CRFs) (GCRFs), remote sensing.**

I. Introduction

AEROSOLS have been recognized among the most important quantities in understanding the Earth's climate [1]. Consequently, the accurate retrieval of aerosol optical depth (AOD or τ), a measure of an extinction of solar radiation by scattering and absorption between the top of the atmosphere and the surface, is of great importance to characterize their effect on the Earth's radiation budget [2]. Currently, a number of satellite-borne sensors monitor the Earth's atmosphere and report their AOD measurements on a daily basis, such as the Moderate Resolution Imaging Spectroradiometer (MODIS) aboard Terra and Aqua satellites [3], Multi-angle Imaging SpectroRadiometer (MISR) aboard Terra [4], Ozone Monitoring Instrument (OMI) aboard Aura [5], or Sea-viewing Wide Field-of-view Sensor (SeaWiFS) aboard SeaStar [6].

The coverage and quality of AOD retrievals from different instruments can vary for a number of reasons. For example, the swath of MODIS is 2330 km, allowing MODIS to cover the entire Earth's surface every day, as opposed to the 360-km swath of MISR local mode, which results in global coverage only every nine days. The quality of AOD estimates from different instruments also varies with atmospheric and surface conditions [7]. In addition to satellite sensors, AOD is measured by ground-based highly accurate sensors from the AErosol RObotic NETwork (AERONET) [8]. AERONET instruments are placed at several hundred unevenly distributed locations across the globe, and their measurements are considered a ground truth. However, AERONET cannot provide the global estimation of AOD required for climate models due to limited spatial coverage.

Different spatial and temporal coverage, design, and specific mission objectives of the instruments mean that they observe and measure different, possibly complementary, aspects of the same phenomenon. Instead of considering various data sources in isolation, combining retrievals from different sources into a unique aggregated AOD retrieval might be the best path toward obtaining a higher quality AOD data product. This was observed in [7], where the simple average of collocated Terra MODIS and MISR retrievals led to an improved accuracy of AOD retrieval. This result indicates that further improvements might be possible if more powerful schemes were used. An equally important issue in the remote sensing of aerosols, in addition to obtaining the point estimate of AOD, is the estimation of retrieval uncertainty. Since AOD retrievals are used as inputs to complex climate models [9], adequate knowledge about the uncertainty and quality of AOD retrievals from satellite instruments is of extreme importance for climate studies.

Several issues need to be considered for the task of aggregation of AOD retrievals. Namely, AOD distribution is characterized by strong temporal and spatial correlation, which could be used to improve the accuracy of aggregated retrievals. Furthermore, due to a number of reasons (e.g., limited coverage of sensors, sensor maintenance, and sunglint), it is common that satellite or ground-based retrievals are missing. In this letter, we propose an aggregation approach that handles these issues, based on a special type of conditional random field (CRF) called Gaussian CRF (GCRF) [10], [11]. The GCRF model can utilize correlations in the values of AOD while allowing learning and inference with missing retrievals. Finally, the approach provides an easy-to-calculate point estimate of AOD, as well as estimation uncertainty.

II. GCRF

GCRF provides a probabilistic framework for the incorporation of various aspects of complex data into a single model. Let us denote a vector of covariates by $\mathbf{x}$ and an N-dimensional vector of real-valued output variables by $\mathbf{y} = [y_1, \ldots, y_N]^{\mathrm{T}}$.

Manuscript received April 7, 2014; revised August 20, 2014; accepted September 29, 2014. Date of publication October 17, 2014; date of current version October 31, 2014. This work was supported by the National Science Foundation under Grant NSF-IIS-1117433.

The authors are with the Department of Computer and Information Sciences, Temple University, Philadelphia, PA 19122 USA (e-mail: nemanja@temple.edu; vladan@temple.edu; zoran@temple.edu; vucetic@temple.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/LGRS.2014.2361154

For example, y_i can be an actual AOD at a particular time and place, while $\mathbf{x}$ are all available measurements related to AOD at different times and locations. The conditional distribution $\mathbb{P}(\mathbf{y}|\mathbf{x})$ for CRF can be represented in a convenient form as

$$\mathbb{P}(\mathbf{y}|\mathbf{x}) = \frac{1}{Z(\mathbf{x},\boldsymbol{\alpha},\boldsymbol{\beta})} \exp\left(-\sum_{i=1}^{N} A(\boldsymbol{\alpha}, y_i, \mathbf{x}) - \sum_{i\sim j} I(\boldsymbol{\beta}, y_i, y_j, \mathbf{x})\right) \tag{1}$$

where $A(\boldsymbol{\alpha}, y_i, \mathbf{x})$ is an *association potential* with weights $\boldsymbol{\alpha}$, $I(\boldsymbol{\beta}, y_i, y_j, \mathbf{x})$ is an *interaction potential* with weights $\boldsymbol{\beta}$, $i \sim j$ denotes that y_i and y_j are assumed correlated (referred to as neighbors), and $Z(\mathbf{x}, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a normalization function.

In general, both learning and inference with the model defined in (1) can be difficult due to integration over real-valued $\mathbf{y}$ in $Z(\mathbf{x}, \boldsymbol{\alpha}, \boldsymbol{\beta})$. However, the potentials could be designed in a way which allows efficient learning and inference. First, let us define the association potential as follows:

$$A(\boldsymbol{\alpha}, y_i, \mathbf{x}) = \sum_{m=1}^{M} \alpha_m \left(y_i - \theta_m\left(\mathbf{x}_i^m\right)\right)^2 \tag{2}$$

where $\theta_m(\cdot)$ is the mth baseline predictor, α_m is the weight of the mth predictor, $\mathbf{x}_i^m$ is a vector of covariates used by the mth predictor to predict y_i, M is the number of prediction models, and $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_M]^{\mathrm{T}}$. Baseline predictor $\theta_m(\cdot)$ can be any predictor of y_i (e.g., operational aerosol retrieval algorithm for a particular instrument). The quadratic function is easy to interpret: The value of y_i close to $\theta_m(\mathbf{x}_i^m)$ is more likely by the model in (1). We can introduce an arbitrary number of baseline predictors, and their relevance will be determined during training: Relevant predictors will be given bigger α weights, whereas irrelevant ones will get weights close to 0, thus reducing their influence.

Furthermore, let us define the interaction potential as follows:

$$I(\boldsymbol{\beta}, y_i, y_j, \mathbf{x}) = \sum_{l=1}^{L} \beta_l \delta_{ij}^l \cdot (y_i - y_j)^2 \tag{3}$$

where L is the number of interaction (or neighborhood) definitions, each assigned a different weight β_l, δ_{ij}^l is a 0/1 indicator function describing whether the ith and the jth outputs are connected according to the lth neighborhood definition, and $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_L]^{\mathrm{T}}$. If two outputs are neighbors (e.g., AOD at the same location for two consecutive days), the interaction potential (3) will force them to have similar values.

When the potentials are defined as in (2) and (3), it can be shown that the resulting CRF model corresponds to a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x}))$ [11]. For this reason, we call the resulting model the GCRF. It is important to observe that both mean and covariance matrix are not constant and that they depend on $\mathbf{x}$. However, for the simplicity of notation, we will use $\boldsymbol{\Sigma} \equiv \boldsymbol{\Sigma}(\mathbf{x})$ and $\boldsymbol{\mu} \equiv \boldsymbol{\mu}(\mathbf{x})$.

To obtain explicit expressions for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, let us first define an N-dimensional vector $\mathbf{b} = [b_1, \ldots, b_N]^{\mathrm{T}}$ with elements

$$b_i = 2 \sum_{m=1}^{M} \alpha_m \theta_m\left(\mathbf{x}_i^m\right) \tag{4}$$

and $N \times N$ matrices $\mathbf{Q}_1$ and $\mathbf{Q}_2$ with elements

$$\mathbf{Q}_{1ij} = \begin{cases} \sum_{m=1}^{M} \alpha_m, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}$$

$$\mathbf{Q}_{2ij} = \begin{cases} \sum_{n=1}^{N} \sum_{l=1}^{L} \beta_l \delta_{in}^l, & \text{if } i = j \\ -\sum_{l=1}^{L} \beta_l \delta_{ij}^l, & \text{if } i \neq j. \end{cases} \tag{5}$$

As shown in [11], the inverse of the covariance matrix (i.e., precision matrix) of $\mathbb{P}(\mathbf{y}|\mathbf{x})$ can be calculated as

$$\boldsymbol{\Sigma}^{-1} = 2(\mathbf{Q}_1 + \mathbf{Q}_2) \tag{6}$$

and the mean of $\mathbb{P}(\mathbf{y}|\mathbf{x})$ can be calculated as $\boldsymbol{\mu} = \boldsymbol{\Sigma}\mathbf{b}$. For a special case where $\beta_1 = \beta_2 = 0$, we can see that the mean of the ith point is equal to a linear combination of $\theta_m(\mathbf{x}_i^m)$ weighted by $\alpha_m / \sum_l \alpha_l$.

A. Training and Inference in the GCRF Model

Given the model from (1) and a training set $\mathcal{D} = (\mathbf{x}, \mathbf{y}) = \{(\mathbf{x}_i, y_i)\}_{i=1,\ldots,N}$, the training task is to find $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that the conditional log-likelihood $\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is maximized

$$(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\beta}}) = \arg\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \text{ where } \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \log \mathbb{P}(\mathbf{y}|\mathbf{x}) \tag{7}$$

solved using gradient descent [11]. On the other hand, given a trained GCRF model with parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, the inference task is to find the point estimate $\hat{\mathbf{y}}$ of outputs $\mathbf{y}$ for given inputs $\mathbf{x}$. We select $\hat{\mathbf{y}}$ that maximizes $\mathbb{P}(\mathbf{y}|\mathbf{x})$, equal to the expected value $\boldsymbol{\mu}$

$$\hat{\mathbf{y}} = \boldsymbol{\mu} = \boldsymbol{\Sigma}\mathbf{b}. \tag{8}$$

An important property of the GCRF model is that the uncertainty of the calculated point estimates is easily computed. More specifically, 95% confidence intervals of outputs $\hat{\mathbf{y}}$ are estimated from the mean and the covariance matrix as

$$\mathbb{P}\left(\hat{\mathbf{y}} - 1.96 \cdot \mathrm{diag}(\boldsymbol{\Sigma}) \leq \mathbf{y} \leq \hat{\mathbf{y}} + 1.96 \cdot \mathrm{diag}(\boldsymbol{\Sigma})\right) = 0.95 \tag{9}$$

where $\mathrm{diag}(\boldsymbol{\Sigma})$ denotes the main diagonal of the $\boldsymbol{\Sigma}$ matrix.

B. Handling the Missing Predictors

In many real-world applications, it is often the case that some baseline predictors $\theta_m(\mathbf{x}_i)$ might not be available. To address this issue, we define the association potential as

$$A(\boldsymbol{\alpha}, y_i, \mathbf{x}) = \sum_{m=1}^{M} \alpha_m \delta_i^m \left(y_i - \theta_m\left(\mathbf{x}_i^m\right)\right)^2 \tag{10}$$

where we introduced 0/1 functions δ_i^m equal to 1 if the mth baseline predictor provided prediction for the ith output and 0 if otherwise. This results in slightly modified expressions (4) and (5) for $\mathbf{b}$ and $\mathbf{Q}_1$, where α_m is replaced by $\alpha_m \delta_i^m$.

III. GCRF for Aggregation of AOD Retrievals

We address the problem of aggregation of satellite AOD retrievals. More formally, we assume that we are given training set $\mathcal{D} = \{\theta_i^{aqua}, \theta_i^{omi}, \theta_i^{sw}, \theta_i^{terra}, \theta_i^{misr}, y_i\}_{i=1,\ldots,N}$, where N is the size of the data set, index i corresponds to AOD at a particular time and location, y_i is an AERONET retrieval taken as a ground truth due to the instrument's high accuracy, and, to make the notation more intuitive, instead of covariates $\mathbf{x}$ and baseline predictors $\theta_m(\mathbf{x}_i^m)$ introduced in (2), $m = 1, \ldots, 5$, we used θ_i^{aqua}, θ_i^{omi}, θ_i^{sw}, θ_i^{terra}, and θ_i^{misr}, respectively, which denote operational AOD retrievals from Aqua MODIS, OMI, SeaWiFS, Terra MODIS, and MISR, respectively. Furthermore,

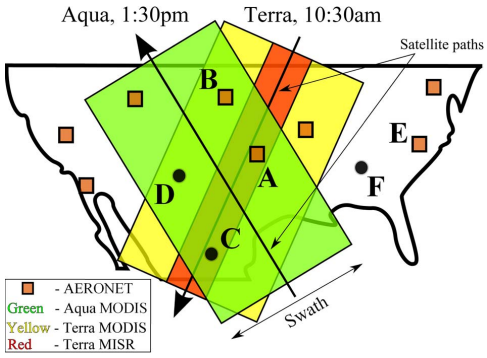

Fig. 1. Coverage of instruments over the USA.

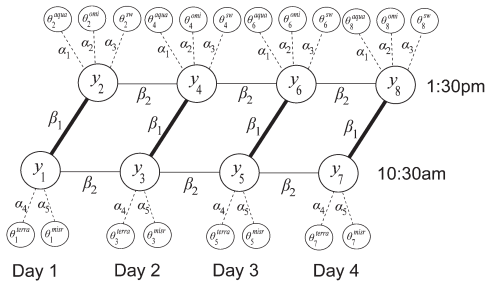

Fig. 2. Graphical representation of GCRF for retrieval aggregation (thick line: within-day interaction; thin line: between-day interaction; dashed line: association between inputs and outputs).

training points with available AERONET retrieval are referred to as labeled; otherwise, they are referred to as unlabeled points.

Due to limitations of the sensors, it is common that $\mathcal{D}$ has a large number of missing satellite and ground-based retrievals. As an example, let us consider a typical spatial coverage of Terra and Aqua MODIS, MISR, and AERONET instruments in the USA during a single day, given in Fig. 1. We can see that different areas may have very different coverages. For most of labeled data points (e.g., data with AERONET retrieval from locations A, B, and E) and unlabeled points (e.g., data without AERONET retrieval from locations C, D, and F), AOD retrievals from some of the satellite instruments are missing. Locations A and C have retrievals from all three satellites, while B and D are outside MISR's swath and do not have its retrieval. Locations E and F are outside the swaths of all instruments and do not have any retrieval available. Moreover, even when a location is covered by an instrument, retrieval availability also depends on the sensitivity of the instruments to technical, atmospheric, and surface conditions, further exacerbating the problem of limited coverage.

Satellite sensors considered in this study can be divided into two groups, one that collocates with AERONET around 10:30 A.M. local time (comprising Terra MODIS and MISR) and another that does so around 1:30 P.M. local time (comprising Aqua MODIS, OMI, and SeaWiFS). We choose to provide aggregated retrievals at these two discrete time points every day at every location. Thus, each location for each day contributes two data points to $\mathcal{D}$, one with AOD retrievals from two satellite instruments with morning overpass and the other with retrievals from three satellites with afternoon overpass.

A. GCRF Aggregation Model

The graphical representation of the GCRF model for AOD retrieval aggregation, derived from the GCRF model from Section II-B, is shown in Fig. 2. We did not consider spatial correlations due to the sparse distribution of AERONET sites in our data set [12], and we set all interaction weights β between ith and jth outputs to zero if i and j correspond to different locations. As a result, different AERONET sites are independent, and Fig. 2 corresponds to a single location. We note that the same $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ parameters are used for all locations. As we consider retrievals from five satellite sensors, we set $M = 5$ in (2) for the association potential and represent the influence of satellite retrievals on outputs with a dashed line in Fig. 2. For the interaction potential, we assume that AOD values are temporally correlated. To encode this assumption, we linked within-day outputs at 10:30 A.M. and 1:30 P.M. and associate weight β_1 with these links (thick lines in Fig. 2). We also linked 10:30 A.M. outputs from two consecutive days, as well as 1:30 P.M. outputs from two consecutive days, and associate weight β_2 with these day-to-day links (thin lines in Fig. 2). The resulting potentials are

$$\begin{aligned} A(\boldsymbol{\alpha}, y_i, \mathbf{x}) =& \alpha_1 \delta_i^1 \left(y_i - \theta_i^{aqua}\right)^2 + \alpha_2 \delta_i^2 \left(y_i - \theta_i^{omi}\right)^2 \\ &+ \alpha_3 \delta_i^3 \left(y_i - \theta_i^{sw}\right)^2 + \alpha_4 \delta_i^4 \left(y_i - \theta_i^{terra}\right)^2 \\ &+ \alpha_5 \delta_i^5 \left(y_i - \theta_i^{misr}\right)^2 \\ I(\boldsymbol{\beta}, y_i, y_j, \mathbf{x}) =& \beta_1 \delta_{ij}^1 \cdot (y_i - y_j)^2 + \beta_2 \delta_{ij}^2 \cdot (y_i - y_j)^2 \end{aligned} \tag{11}$$

where δ_i^m, $m \in \{1, \ldots, 5\}$, are 0/1 indicator functions returning 1 if, for the ith output, there is an available retrieval from Aqua MODIS, OMI, SeaWiFS, Terra MODIS, and MISR, respectively, and 0 if otherwise, while δ_{ij}^1 and δ_{ij}^2 are indicator functions returning 1 if y_i and y_j are within-day neighbors or between-day neighbors, respectively, and 0 if otherwise.

The GCRF aggregation model has very useful properties. In particular, it can aggregate baseline retrievals even when an arbitrary number of baseline predictors are unavailable for some output y_i. GCRF learns the importance of different baselines by assigning higher α weight to more accurate ones and can also utilize temporal correlations in AOD values, where the strength of temporal correlations is learned and quantified through β weights. Moreover, GCRF readily provides uncertainty estimates of output retrievals. Finally, assuming that we use indexing as shown in Fig. 2, the resulting matrix $\boldsymbol{\Sigma}^{-1}$ in (6) is pentadiagonal; thus, its inverse, required in the gradient-based optimization, can be found in very favorable $\mathcal{O}(KN)$ time [13]. Note that, in general, gradient may be more expensive to compute, depending on a specific GCRF graph structure.

B. GCRF Training With Missing AERONET AOD Retrievals

As discussed previously, it is possible to have unlabeled data points in the training data. In that case, the aggregation becomes a semisupervised task [14]. In particular, let us denote the joint probability of labeled and unlabeled points as $\mathbb{P}(\mathbf{y}_L, \mathbf{y}_U | \mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and let us separate prediction vector $\boldsymbol{\mu}$ and precision matrix $\boldsymbol{\Sigma}^{-1}$ into labeled and unlabeled parts as $\boldsymbol{\mu} = [\boldsymbol{\mu}_L^{\mathrm{T}}, \boldsymbol{\mu}_U^{\mathrm{T}}]^{\mathrm{T}}$ and $\boldsymbol{\Sigma}^{-1} = [\mathbf{Q}_{LL}, \mathbf{Q}_{LU}; \mathbf{Q}_{UL}, \mathbf{Q}_{UU}]$. Then, the marginal distribution of the labeled part is equal to $\mathbb{P}(\mathbf{y}_L | \mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}_L, (\mathbf{Q}_{LL} - \mathbf{Q}_{LU}\mathbf{Q}_{UU}^{-1}\mathbf{Q}_{UL})^{-1})$. To obtain a semisupervised training procedure, we redefine the likelihood as $\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \log \mathbb{P}(\mathbf{y}_L | \mathbf{x})$ and maximize the modified (7).

TABLE I
COVERAGE OF BASELINE PREDICTORS IN THE FIVE-YEAR DATA

Instrument	No. of days	Coverage
Terra MODIS	20,286	20.37%
MISR	4,906	4.93%
Aqua MODIS	16,711	16.78%
OMI	23,332	23.43%
SeaWiFS	4,567	4.59%

IV. EXPERIMENTS

We used ground-based AERONET data [8] and data from five satellite instruments from 2006 to 2010, collected over North America, and considered AOD at 550-nm wavelength. If a source did not provide AOD retrievals at this wavelength, we performed a linear interpolation or extrapolation in the log scale of retrievals at the two closest wavelengths to 550 nm [3]. There were, on average, 56 working AERONET sites each year, and we estimated AOD twice a day at every AERONET location, resulting in a total of 199 134 data points in the collected data set. In all experiments, we report results after leave-one-year-out cross-validation. After five repetitions, test data were pooled together and root-mean-squared error (RMSE) on labeled test points was calculated and reported.

For AERONET, we downloaded data from the AERONET website[1] [8] and used the highest quality Level 2.0 AOD data. To obtain satellite data, we used Multi-sensor Aerosol Products Sampling System (MAPSS)[2] [15]. In addition to AOD, MAPSS contains confidence assigned to each retrieval (quality assurance (QA) flags). Only highest quality retrievals were retained in order to reduce biases inherent to individual instruments, as suggested in [16]. We used the following AOD data products: 1) **MODIS**—Daily Level 2 product, collection 5.1 (MOD04_L2 and MYD04_L2 for Terra and Aqua, respectively); 2) **MISR**—MIL2ASAE, a MISR Level 2 product; 3) **OMI**—OMAERUV, a Level-2 near-UV product; and 4) **SeaWiFS**—SWDB_L2, Deep Blue Daily Level 2 product. For each satellite source, all retrievals within 30 km of an AERONET site were collected and averaged. In order to further reduce biases, Petrenko and Ichoku [16] also suggest outlier removal based on a z-score test that requires ground-truth data. This prevents the removal of outliers in locations without such labels, and as the proposed method aims at global applicability, the outlier removal was not applied. Coverage of operational retrievals is given in Table I. We see that MISR and SeaWiFS cover less than 5% of days during the five-year period, while OMI has the coverage of nearly 24%.

We report RMSE on subsets of labeled points, shown in Table II. For example, row "$MODIS + OMI$" reports results on a subset of labeled points that have both Aqua MODIS and OMI retrieval, but not SeaWiFS retrieval. We do not list RMSE when a model did not have a complete coverage on test data.

Let us first discuss results of baseline predictors shown in the "*Individual sensors*" column. The results indicate that the performance of different instruments varied significantly, both in accuracy and coverage. Regarding the 1:30 P.M. results, we see that SeaWiFS retrievals overall were more accurate than both Aqua MODIS and OMI retrievals. It is interesting to observe that OMI accuracy was particularly low when MODIS and SeaWiFS retrievals were not available. Whenever OMI provided AOD retrieval along with some other instrument, its accuracy improved significantly, which indicates that there could be certain issues with the quality checks of OMI retrievals. We can also see that SeaWiFS consistently outperformed Aqua MODIS but that it had around three times smaller coverage. Regarding 10:30 A.M. results, MISR achieves 38% lower RMSE than Terra MODIS while having four times smaller coverage. Interestingly, the accuracy of MODIS in the absence of MISR was much higher than its accuracy when MISR was available, which is explained by larger MODIS sensitivity to sunglint [4].

Furthermore, we considered GCRF models with increasing levels of complexity and investigated how the introduction of β parameters influences performance. For that purpose, we trained GCRF with different combinations of β parameters: 1) model without interactions ($\beta_1 = \beta_2 = 0$); 2) model with diurnal interaction ($\beta_1 \neq 0$ and $\beta_2 = 0$); 3) model with day-to-day interaction ($\beta_1 = 0$ and $\beta_2 \neq 0$); and 4) model with both interactions ($\beta_1 \neq 0$ and $\beta_2 \neq 0$). In the bottom three rows, we see that the overall RMSE on all labeled points dropped significantly when we included more temporal interactions in the model. Overall, RMSE on labeled points with at least one satellite retrieval dropped from 0.1430 to 0.0852 when both day-to-day and diurnal interactions were included, which is an improvement of 41%.

Compared to the RMSE of individual instruments, GCRF accuracies for 1:30 P.M. have improved. For 10:30 am, RMSE on points with MISR retrievals increased due to temporal averaging enforced by β_1 and β_2. Thus, it seems that we should use retrievals from very accurate MISR retrievals when available and use the power of temporal smoothing when MISR is missing. However, it is important to note that only 4.76% of data points had high-quality MISR or SeaWiFS retrievals. Moreover, results in the upper and middle parts of Table II cover only 22 420 labeled points that had retrievals from at least one of the sensors. This amounts to 11.2% of the data set, while GCRF provides aggregated retrieval for all 199 134 points due to its ability to leverage temporal correlation.

In Fig. 3, we give an example of daily uncertainty estimates for the AERONET site at the Maryland Science Center in Baltimore, USA. Colored circles represent the availability of individual satellite retrievals. Uncertainty gradually increases with distance from the nearest retrieval, which is an example of the influence of the β_2 parameter modeling daily interactions. It is interesting to observe a drop in uncertainty at day 185 in Fig. 3(b), although there was no instrument-retrieved AOD at 1:30 P.M. on that or any of the neighboring days. This is due to the Terra MODIS retrieval observed at 10:30 A.M. that day, as seen in Fig. 3(a). This exemplifies the influence of β_1, which models the interaction between 10:30 A.M. and 1:30 P.M. retrievals. The influence of $\boldsymbol{\alpha}$ parameters is also visible in Fig. 3; the availability of MISR or SeaWiFS retrievals results in sharp uncertainty drops as these instruments were assigned the largest α parameters.

V. CONCLUSION

We presented a GCRF model for fusion of AOD retrievals from multiple instruments. The ease of modeling interaction between outputs, ability to handle missing data, high quality

[1] aeronet.gsfc.nasa.gov/cgi-bin/combined_data_access_new, March 2014
[2] disc.sci.gsfc.nasa.gov/aerosols/services/mapss/mapssdoc, March 2014

TABLE II
RMSE OF INSTRUMENTS AND GCRF (UPPER PART SHOWS 1:30 P.M., MIDDLE PART SHOWS 10:30 A.M., AND BOTTOM PART SHOWS LABELED DATA RESULTS)

Coverage	No. of points	Individual sensors	Only β_1	Only β_2	All β
All MODIS Aqua	7,246	0.0872	0.0822	0.0825	**0.0756**
All OMI	10,142	0.2390	0.1930	0.1482	**0.0934**
All SeaWiFS	2,205	0.0739	0.0630	0.0642	**0.0607**
MODIS Aqua alone	2,102	0.0889	0.0850	0.0773	**0.0741**
OMI alone	4,528	0.2934	0.2744	0.2010	**0.1114**
SeaWiFS alone	237	0.0800	0.0771	0.0784	**0.0753**
MODIS + OMI	3,868	0.0893 | 0.2123	0.0886	0.0925	**0.0827**
MODIS + SeaWiFS	222	0.0982 | 0.0747	0.0665	0.0655	**0.0634**
OMI + SeaWiFS	692	0.1011 | 0.0837	0.0758	0.0760	**0.0691**
MODIS + OMI + SeaWiFS	1,054	0.0717 | 0.0715 | 0.0650	**0.0475**	0.0506	0.0495
All MODIS Terra	8,725	0.0905	0.0826	0.0819	**0.0800**
All MISR	2,165	**0.0652**	0.0664	0.0735	0.0725
MODIS Terra alone	7,552	0.0863	0.0845	0.0822	**0.0806**
MISR alone	992	**0.0618**	0.0636	0.0657	0.0677
MODIS + MISR	1,173	0.1142 | **0.0680**	0.0686	0.0795	0.0763
All labeled	44,445	–	0.1634	0.1328	**0.1080**
All labeled with any satellite	22,420	–	0.1430	0.1158	**0.0852**
All labeled without satellites	20,025	–	0.1817	0.1481	**0.1271**

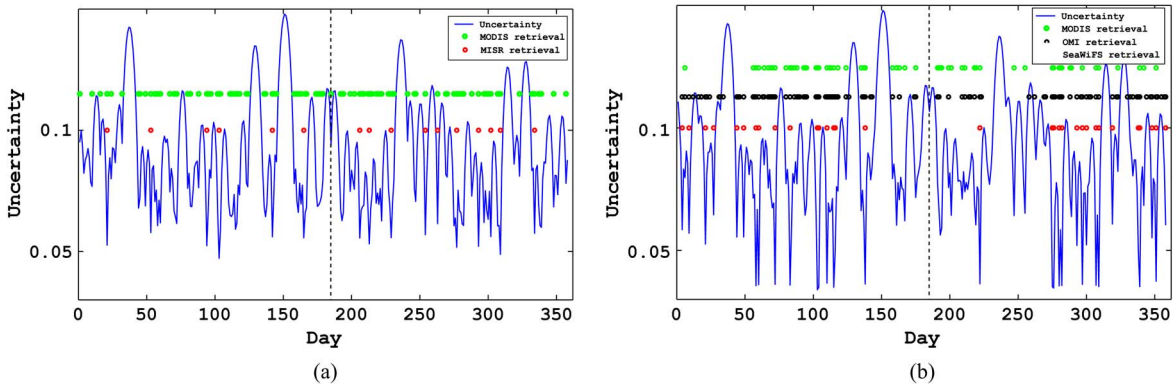

Fig. 3. GCRF uncertainty estimate for AERONET site at Maryland Science Center for 2007 (dashed line denotes day 185).

of the aggregated AOD retrievals, and the interpretability of its outputs strongly suggest that the GCRF model can represent an important tool in remote sensing applications.

ACKNOWLEDGMENT

The authors would like thank the AERONET Principal Investigators and their staff for establishing and maintaining sites used in this study.

REFERENCES

[1] R. J. Charlson *et al.*, "Climate forcing by anthropogenic aerosols," *Science*, vol. 255, no. 5043, pp. 423–430, Jan. 1992.
[2] R. C. Levy, L. A. Remer, and O. Dubovik, "Global aerosol optical properties and application to moderate resolution imaging spectroradiometer aerosol retrieval over land," *J. Geophys. Res.*, vol. 112, no. D13, pp. 13 210–13 224, Jul. 2007.
[3] L. A. Remer *et al.*, "The MODIS aerosol algorithm, products, validation," *J. Atmos. Sci.*, vol. 62, no. 4, pp. 947–973, Apr. 2005.
[4] R. A. Kahn *et al.*, "MISR aerosol product attributes and statistical comparisons with MODIS," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 12, pp. 4095–4114, Dec. 2009.
[5] O. Torres, R. Decae, J. P. Veefkind, and G. de Leeuw, "OMI Aerosol Retrieval Algorithm," in *OMI Algorithm Theoretical Basis Document, Volume III, Clouds, Aerosols, Surface UV Irradiance*, 2002.
[6] M. Wang, S. Bailey, and C. R. Mcclain, "SeaWiFS provides unique global aerosol optical property data," *EOS, Trans. Amer. Geophys. Union*, vol. 81, no. 18, pp. 197–202, May 2000.
[7] M. I. Mishchenko *et al.*, "Toward unified satellite climatology of aerosol properties. 3. MODIS versus MISR versus AERONET," *J. Quant. Spectr. Radiat. Transf.*, vol. 111, no. 4, pp. 540–552, 2010.
[8] B. N. Holben *et al.*, "AERONET—A federated instrument network and data archive for aerosol characterization," *Remote Sens. Environ.*, vol. 66, no. 1, pp. 1–16, Oct. 1998.
[9] K. Zhang *et al.*, "The global aerosol-climate model ECHAM-HAM version 2: Sensitivity to improvements in process representations," *Atmos. Chem. Phys.*, vol. 12, no. 19, pp. 8911–8949, 2012.
[10] T. Qin, T. Yan Liu, X. dong Zhang, D. sheng Wang, and H. Li, "Global ranking using continuous conditional random fields," in *Proc. Adv. Neural Inf. Process. Syst. 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds., 2008, pp. 1281–1288.
[11] V. Radosavljevic, S. Vucetic, and Z. Obradovic, "Continuous conditional random fields for regression in remote sensing," in *Proc. ECAI*, 2010, pp. 809–814.
[12] M. Masmoudi *et al.*, "Spatial and temporal variability of aerosol: Size distribution and optical properties," *Atmos. Res.*, vol. 66, no. 1/2, pp. 1–19, Mar. 2003.
[13] A. D. A. Hadj and M. Elouafi, "A fast numerical algorithm for the inverse of a tridiagonal and pentadiagonal matrix," *Appl. Math. Comput.*, vol. 202, no. 2, pp. 441–445, Aug. 2008.
[14] X. Zhu, Z. Ghahramani, and J. Lafferty, "Semi-supervised learning using Gaussian fields and harmonic functions," in *Proc. ICML*, 2003, pp. 912–919.
[15] M. Petrenko, C. Ichoku, and G. Leptoukh, "Multi-Sensor Aerosol Products Sampling System (MAPSS)," *Atmos. Meas. Techn.*, vol. 5, no. 5, pp. 913–926, 2012.
[16] M. Petrenko and C. Ichoku, "Coherent uncertainty analysis of aerosol measurements from multiple satellite sensors," *Atmos. Chem. Phys.*, vol. 13, no. 14, pp. 6777–6805, 2013.