A New Version of the SUNY Solar Forecast Model: A Scalable Approach to Site-Specific Model Training

Richard Perez[1], James Schlemmer[1], Sergey Kivalov[1], John Dise[2],
Patrick Keelin[2], Mark Grammatico[2], Thomas Hoff[2], & Aidan Tuohy[3]

[1] Atmospheric Sciences Research Center, SUNY, Albany, New York, 12203, USA
[2] Clean Power Research, Napa, California, 94558, USA
[3] Electric Power Research Institute, Palo Alto, California, 94304, USA

Abstract — **This article introduces a new version of the SUNY solar forecast model, as implemented in the software SolarAnywhere®. Like the existing version, this new version is intended for direct, out-of-the-box application throughout North America without requiring training/feedback from measured data. The existing version was recently identified by EPRI as most accurate among thirteen operational models after an independent evaluation in two climatically distinct US regions. This new version shows further measurable performance improvements and operational functionality with capability of using historical satellite data for model training purposes. The advantage of incorporating historical satellite data is that it allows this forecast application to scale to all sizes of PV systems (large utility-scale down to individual rooftop) without the requirement of site-measured inputs. This new approach exhibits a 1-5% root mean square error (RMSE) improvement over the forecast horizon up to two-days in advance.**

Index Terms —**solar resource, irradiance, solar forecasting, PV fleet forecasting.**

I. INTRODUCTION

The existing operational SUNY forecast model (Version 4, or V4) is based upon an optimized mix of models including a satellite-based cloud motion vector (CMV) model, and global/regional solar and cloud cover Numerical Weather Prediction (NWP) models, [1].

Coupled with a PV simulation engine (e.g., [2]), V4 is readily deployable on any geographical scale, from single plants to dispersed PV fleets, without requiring system data training. It is applicable anywhere the underlying NWP and cloud motion models are available – i.e., potentially anywhere on the planet. Current operational coverage is North America, with planned expansions in Asia and South America.

Out-of-the-box application without localized training is a key attribute for distributed fleets where PV system specs exist (e.g., from [3]) but where measured quality training data are seldom available. However, it is important to stress that this key attribute does not preclude model in-situ training if and when possible.

V4 was recently evaluated independently by EPRI at test locations in North Carolina and California [4]. It was found to be the most accurate operational model among thirteen US-based models from various providers – noting that several of the models evaluated where trained (e.g., via deep machine learning) using measured plants data made available by EPRI to forecast providers.

Here we present and evaluate a new SolarAnywhere version (V4x) that retains the same philosophy (i.e., readily deployable without system data training) but pushes operational accuracy further with a better exploitation of operational model input data, and utilization of SolarAnywhere historical data [5].
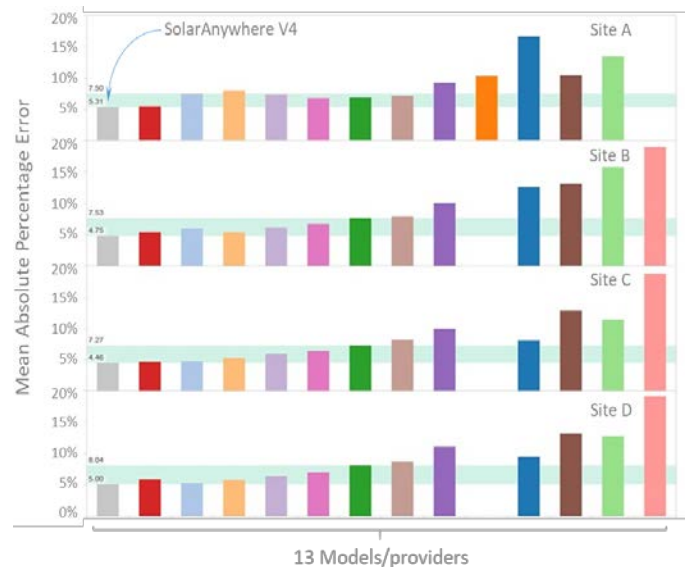

Fig. 1. Results of the EPRI forecast evaluation trial

II. MODEL ENHANCEMENTS

Site-Independent Enhancement: The existing operational V4 is derived from an optimized blend of a CMV model, and four NWP models – ECMWF, GFS, NDFD and HRRR [1]. The blend is a function of time horizon and time of day. The new version adds predicted solar conditions as an additional optimization input. The new optimum blend is determined empirically from 10 months of measured data at seven SURFRAD network locations [6]. These locations sample a diversity of CONUS climatic environments.

Location-Specific Improvements: In a recent publication [7], we observed that historical satellite-derived irradiances were as effective as ground measurements to benchmark forecast models, producing comparable error metrics across climatically distinct locations and forecast time horizons (Figure 2). We take advantage of this observation as an opportunity to regionalize input blend optimization and, as needed, remove condition-specific biases. Historical SolarAnywhere data are available for long time spans over entire continents and can thus be applied to locally optimize forecasts blends and biases in any region/location. Of course, this assumes that historical SolarAnywhere data are accurate enough and fully representative of the considered region/location. A growing body of evidence indicates that this is the case (e.g., [8]).
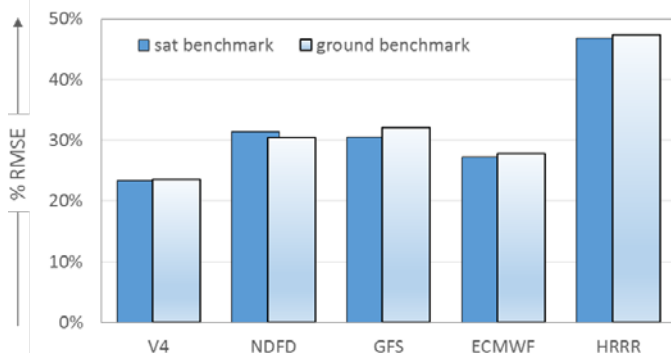

Fig. 2. Comparing satellite and ground benchmarked forecast RMSE statistics across all time horizons at multiple locations [7]

Probabilistic functionality: In addition to deterministic performance improvements, V4.x also produces probabilistic information. The benefit of historical satellite irradiance data is the applicability for probabilistic forecasting without the requirement of ground-based measurement. Each deterministic forecast produces a probabilistic envelope from which operational probability quantiles can readily be extracted. The probabilistic envelopes are *experience based*, i.e., they are derived empirically from multi-site validations using either ground measurements or historical satellite data -- we use the same 10-month, seven-SURFRAD site sample as in the above blend optimization to derive the probabilistic envelopes.

Experimental Data: The data assembled for this study spans a period from July 2015 to April 2016 and includes:

- Hourly Ground measurements for the SURFRAD sites [6] of Bondville, IL, Boulder, CO, Desert Rock, NV, Fort Peck MT, Goodwin Creek, MS, Penn State, PA and Sioux Falls, SD.
- Historical intermediate-resolution SolarAnywhere Version 3 (SAV3) satellite-derived irradiances at arbitrary location in North America (hourly, 10 kilometer resolution) [8]
- Satellite-derived cloud motion vector (CMV) forecasts at arbitrary location in North America with 1-5 hour time horizons [1]
- ECMWF, GFS, NDFD and HRRR NWP forecasts at arbitrary locations in North America with 1-48 hour time horizons [1], except for HRRR (1-16 hours only).

III. RESULTS

In figure 3, we report relative RMSEs for all models as a function of time horizons (1-48 hours ahead) across all measurement locations. The models include previous, current and new versions of the SUNY/SolarAnywhere models as well as their underlying model components, smart persistence. The historical satellite model is also plotted as a visual baseline reference. Smart persistence follows the definition of IEA Task 36 [10], i.e., not only accounting for solar geometry detrending, but also increasing the integration time of the reference measurement commensurately with forecast horizon. This definition is considerably more stringent (hence model taxing) than the commonly used, and often mislabeled, smart-persistence that accounts for solar geometry-only.

Existing SUNY/SolarAnywhere forecast model versions include the original V2.4 [1] and the currently operational V4 model. The new versions include a site-independent model (V4x), a location-specific model locally trained with ground measurements (V4x site-specific), and a location-specific model trained with satellite-derived irradiances (V4x site-specific-satellite). In Figure 4, we compare the skill of all versions at the operationally important 3-hour and 24-hour forecast time horizons. Skill is calculated in reference to the IEA smart persistence [10].

Table one reports relative RMSE and MAE for individual SURFRAD locations for one, 3 and 24 hours forecast horizons.

Site-independent vs. location-specific V4x forecasts: The new site-independent model (V4x) shows systematic performance improvement at all locations and time horizons compared to the current operational model (V4). About half of the gain in performance originates from a better model blending fit to the ensemble of the ground locations (V4 had been fitted
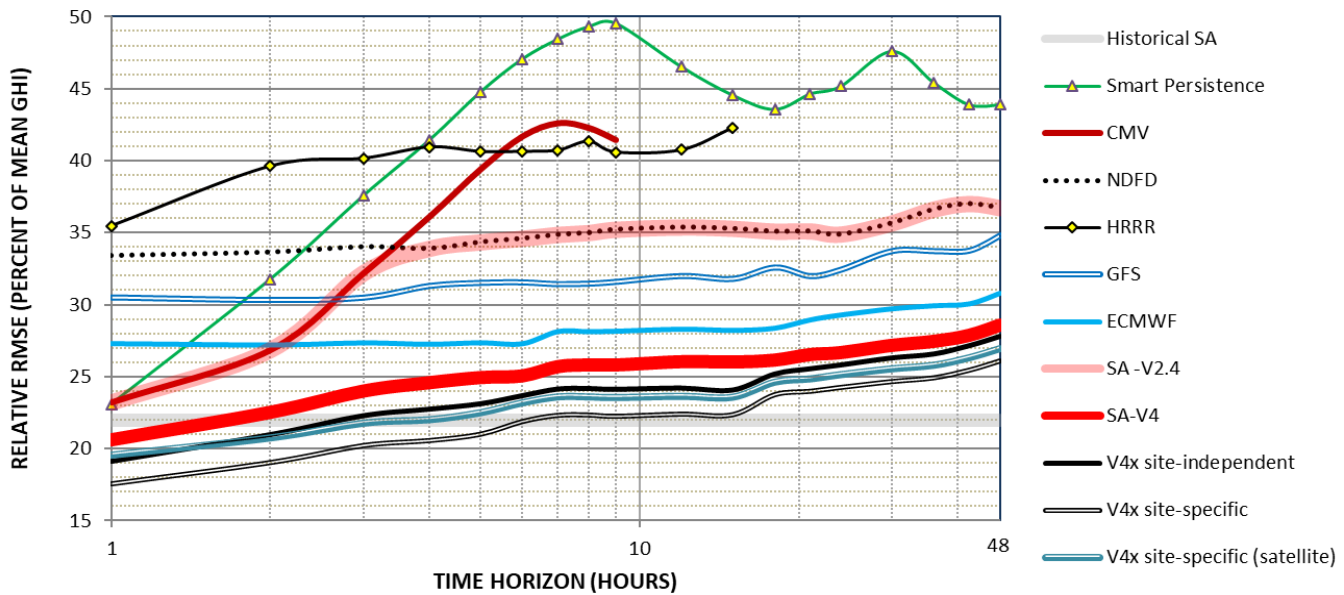
Fig. 3. Relative forecast RMSE as a function of time horizon for the ensemble of seven SURFRAD locations.

to a three-location subset). The other half in performance gain derives from adding predicted insolation conditions as an input to optimum model blending. This improved site-independent model fully retains the "out-of-the-box" applicability of the V4 version tested by EPRI (see Figure 1).

The location-specific version exhibits further performance improvement. Forecast skill exceeds 45% across all time horizons. For this site-specific model, much of the gain in performance is attributable to its parameterization that accounts for predicted insolation conditions. In Figure 5, we compare the performance improvement relative to V4 of different model configurations: The light blue bars correspond to the site-independent model and the darker blue bars correspond to site-fitted models. The bars at left represent model versions that are
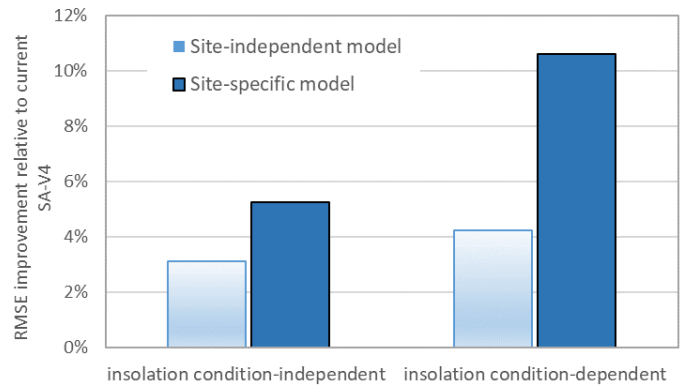

Fig. 5. Comparing the impact of insolation condition parameterization for site-independent and location-specific models
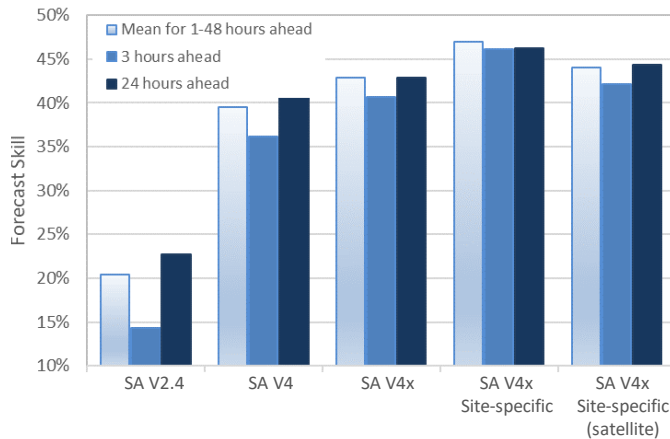
identical to V4 but respectively fitted to all, and individual SURFRAD sites. The bars at right represent V4x models that optimize underlying model blend as a function of predicted insolation conditions. While insolation conditions parameterization increases performance steadily for the site-independent model, the impact is considerably stronger for the site-specific model where the model blend reflects the conditions-dependent strengths and weaknesses of its underlying CMV and NWP components for specific locations.

Operationally, the location-dependent forecast model performance represents a model that would be optimized locally from e.g., measured irradiances or PV plant-output data. This localized performance improvement could plausibly be further improved with real-time feedback techniques -- e.g., using DML or other techniques not addressed in this paper.

Fig. 4. Forecast skill benchmarked to IEA smart persistence

			Bondville	Boulder	Desert Rock	Fort Peck	Goodwin Creek	Penn State	Sioux Falls	All sites
Relative RMSE (% of mean GHI)	1 hour ahead	SA-V2.4	20.2%	23.2%	15.8%	21.1%	17.3%	24.9%	24.1%	20.9%
		SA-V4	20.5%	22.6%	14.9%	20.9%	18.3%	24.9%	22.2%	20.6%
		SA-V4x	19.2%	21.4%	14.1%	19.1%	16.8%	23.3%	20.0%	19.1%
		SA-V4x site spec.	17.8%	20.0%	12.7%	17.2%	15.4%	21.7%	18.3%	17.6%
		SA-V4x site spec. sat.	19.7%	21.3%	15.0%	19.0%	17.6%	23.7%	20.5%	19.5%
	3 hour ahead	SA-V2.4	30.6%	32.1%	24.8%	29.2%	27.8%	34.3%	32.9%	30.2%
		SA-V4	26.4%	26.2%	18.1%	25.3%	25.3%	31.6%	25.1%	25.4%
		SA-V4x	24.9%	24.3%	16.8%	23.9%	22.8%	28.9%	23.7%	23.6%
		SA-V4x site spec.	22.7%	22.6%	14.7%	21.3%	21.0%	26.7%	21.1%	21.4%
		SA-V4x site spec. sat.	24.2%	23.8%	16.8%	23.1%	22.3%	28.1%	23.0%	23.1%
	24 hour ahead	SA-V2.4	35.5%	36.3%	24.6%	35.6%	33.1%	43.0%	37.7%	35.1%
		SA-V4	30.0%	29.5%	20.6%	31.3%	28.2%	33.9%	29.5%	29.0%
		SA-V4x	29.3%	28.5%	20.1%	29.6%	27.2%	33.0%	28.7%	28.1%
		SA-V4x site spec.	27.8%	27.1%	18.3%	27.5%	25.8%	31.4%	26.8%	26.4%
		SA-V4x site spec. sat.	28.8%	27.8%	19.6%	28.4%	26.6%	32.5%	27.7%	27.3%
Relative MAE (% of mean GHI)	1 hour ahead	SA-V2.4	15.8%	16.0%	9.7%	16.2%	13.7%	18.3%	17.6%	15.3%
		SA-V4	14.6%	14.5%	9.1%	14.9%	12.5%	17.5%	15.0%	14.0%
		SA-V4x	13.3%	13.5%	8.2%	13.1%	11.1%	15.5%	13.4%	12.6%
		SA-V4x site spec.	12.1%	12.7%	7.6%	11.9%	10.2%	14.7%	12.1%	11.6%
		SA-V4x site spec. sat.	13.7%	13.3%	8.6%	13.2%	11.8%	16.2%	13.9%	12.9%
	3 hour ahead	SA-V2.4	22.0%	21.7%	14.7%	20.9%	19.9%	24.3%	23.0%	20.9%
		SA-V4	18.4%	17.8%	10.9%	17.5%	17.4%	22.1%	16.6%	17.3%
		SA-V4x	16.7%	15.5%	9.7%	15.6%	14.9%	19.5%	15.4%	15.3%
		SA-V4x site spec.	15.2%	14.4%	8.6%	14.5%	13.4%	18.2%	13.8%	14.0%
		SA-V4x site spec. sat.	16.6%	15.1%	9.8%	15.6%	14.7%	19.2%	15.7%	15.2%
	24 hour ahead	SA-V2.4	24.5%	24.9%	14.7%	24.6%	21.7%	29.2%	25.0%	23.5%
		SA-V4	20.5%	19.4%	12.0%	20.7%	18.6%	23.5%	19.0%	19.1%
		SA-V4x	20.0%	18.6%	12.0%	18.8%	17.9%	22.7%	18.6%	18.4%
		SA-V4x site spec.	19.0%	17.5%	10.7%	18.2%	16.3%	21.4%	17.6%	17.2%
		SA-V4x site spec. sat.	20.6%	17.7%	11.1%	18.2%	17.4%	22.1%	18.5%	17.9%

Satellite-derived site-specific model: The key question here is how much of the site-specific performance enhancement discussed above can be captured when one does not have access to local measurements. SolarAnywhere historical data represent a robust, validated proxy for local measurements.

We find that using SolarAnywhere historical data to locally optimize model configuration does capture some, but not all the localized performance improvement potential. As shown in Figure 4, while the skill of the SolarAnywhere-trained site-specific model is markedly higher than both the existing V4 model and the new site-independent V4x model, it does not reach the performance level of the ground-derived site-specific model -- particularly for shorter time horizons. This is likely because the CMV model is a direct by-product of the historical satellite model used for optimization, which could bias model blending for short horizons in favor of CMV. Nevertheless, the systematic site-specific improvement achievable from the use of SolarAnywhere historical data (available locally throughout entire continents) is significant.

The scatter plots in figure 6 qualitatively illustrate the performance of the current SolarAnywhere V4 model vs the two site-specific versions of the new V4x model: ground-trained, and satellite-trained.

Application to regional forecast model optimization: We provide an example of site/regional model optimization for the State of California. This regional optimization work is undertaken as part of CEC's EPIC research program [9] to develop operational forecast for the climatic regions in California (Figure 6).

This investigation builds on the observations (1) that historical satellite-irradiances are as effective as measurements for model validation (see figure 2), and (2) that satellite-based forecast training is a relatively effective proxy for ground-based training as discussed above.
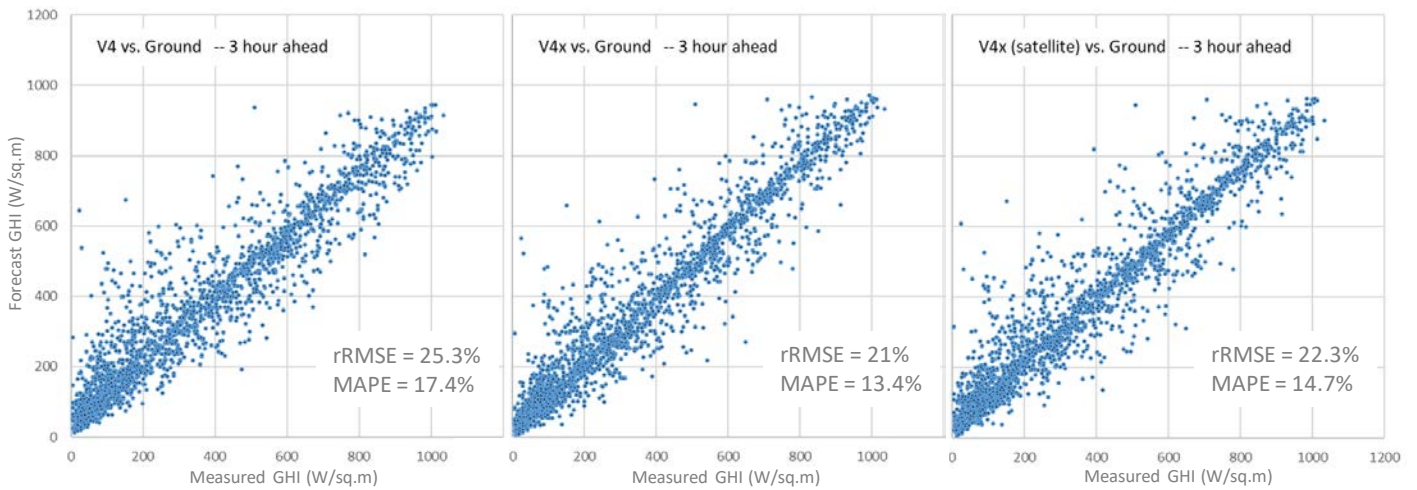
Fig. 5. Comparing performance of the current SolarAnywhere V4 model (left), the site-specific ground-trained V4x model (center) and the satellite-trained V4x model (right) for 3-hours ahead forecasts in Goodwin Creek, MS
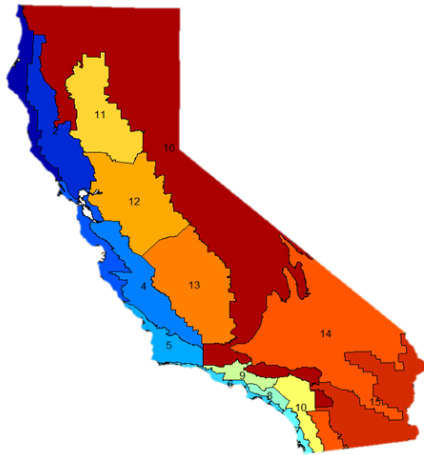

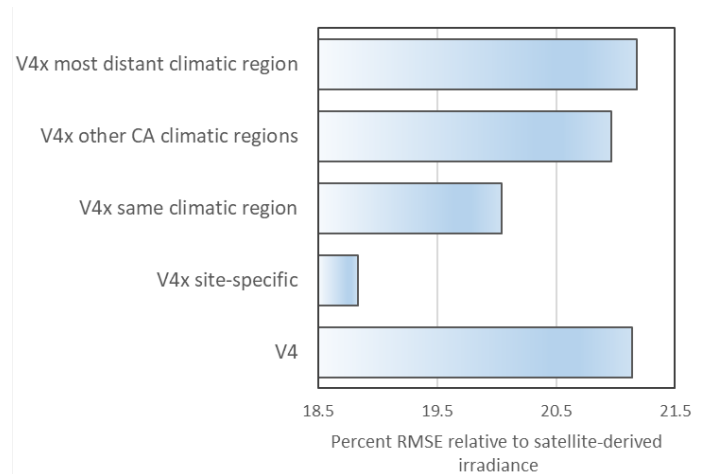Fig. 6. CEC Climatic regions of California

Fig. 7. Resulting percent RMSEs for existing and locally fitted versions of the model across all forecast horizons (1-48 hours) for four California climatic regions.

Here we present preliminary results for four climatically distinct California regions:

- Region 1-- Northern coastal strip
- Region 9 -- Los Angeles Hills
- Region 11 -- Northern Sacramento Valley
- Region 15 – Imperial Valley Desert

Figure 7 contrasts the (relative RMSE-benchmarked) performance of the current and new model versions across a sample of locations in these four regions. Relative to the current V4 version, the locally optimized model exhibits a performance improvement comparable to the site-specific results presented above. Models fitted from other intra-region points retain some of the performance gain. However, models optimized with other regions' points do not show improvement and may show some performance degradation in the worst case.

Probabilistic functionality: As a preliminary example of the new capabilities of the SUNY forecast engine – delivering both deterministic and probabilistic information -- Figure 5 shows probabilistic forecast quantiles as a function of time horizon for a predicted irradiance equal to 500 Wm^{-2}. The figure compares probabilistic quantiles using site-independent models derived from ground data (the seven SURFRAD sites sample) and from historical satellite data at the same locations.

IV. CONCLUSIONS

We presented three versions of the operational SolarAnywhere forecast engine: a site-independent model readily deployable anywhere in North America, a locally

trained model from historical ground measurements, and a locally trained model from historical satellite-derived measurements. All three versions show measureable performance improvement over the current version that had already been singled out as best performer among thirteen US forecast services.

Finally, we showed that model training using SolarAnywhere irradiances could capture a substantial fraction of performance improvement achievable with ground measurement training.
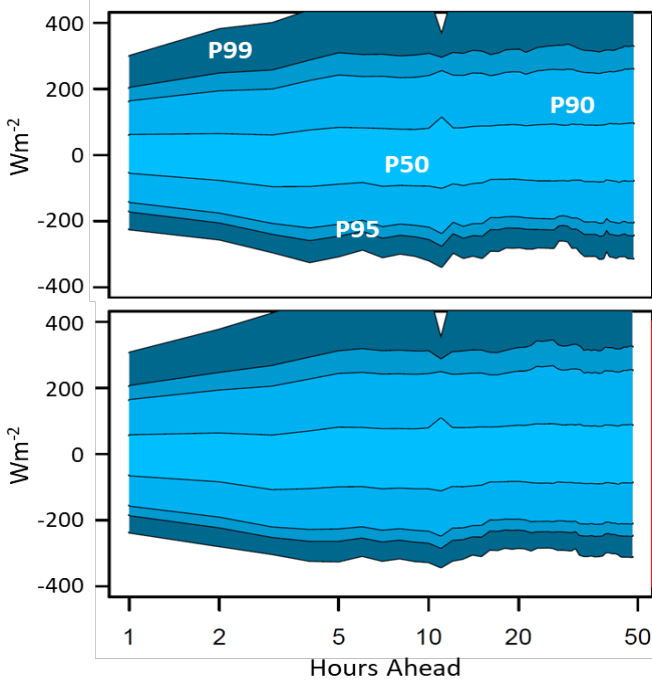

Fig. 5. Experienced-based modeled forecast probability quantiles as a function of time horizon for a predicted irradiance of 500 Wm-2. The top graph illustrates a model empirically derived from ground measurements. The bottom graph illustrates a model empirically derived from SolarAnywhere historical irradiances

REFFENCES

[1] Perez, R., A. Kankiewicz, J. Schlemmer, K. Hemker , Jr., and S. Kivalov, (2014): "A New Operational Solar Resource Forecast Service for PV Fleet Simulation". in 40th IEEE PV Specialists Conference, 2014

[2] For instance, SolarAnywhere® FleetView® https://www.solaranywhere.com/products/solaranywhere-fleetview/

[3] e.g., https://www.cleanpower.com/products/powerclerk/

[4] EPRI (2107): Solar Power Forecasting for Grid Operations: Evaluation of Commercial Providers. EPRI Technology Insights, Nov. 2017.

[5] https://www.solaranywhere.com/

[6] NOAA Earth System Research Laboratory, Global Monitoring Division, (2016): Notice of SURFRAD Network Data Problem Reports. https://www.esrl.noaa.gov/gmd/grad/surfrad/problems.html

[7] Perez, R., J. Schlemmer, K. Hemker, Jr., S. Kivalov, A. Kankiewicz and J. Dise, (2016): "Solar Energy Forecast Validation for Extended Areas & Economic Impact of Forecast Accuracy". In 43th IEEE PV Specialists Conference, 2016

[8] Perez, R., J. Schlemmer, A. Kankiewicz, J. Dise, A. Tadese, & T. Hoff, (2017): Detecting Calibration Drift at Ground Truth Stations. A Demonstration of Satellite Irradiance Models' Accuracy. In 43th IEEE PV Specialists Conference, 2017

[9] California Energy Commission (2017): Solar +: Taking the Next Steps to Enable Solar as a Distribution Asset. Group 5: Holistic Forecasting to Support High-Penetration Solar Grid Operations.

[10] IEA SHC Task 36 (2011): International Energy Agency, Solar Heating and Cooling Programme, Task 36 Solar Resource Knowledge Management, subtask A, model benchmarking http://www.iea-shc.org/task36/