CALIFORNIA ENERGY COMMISSION

Energy Research and Development Division

FINAL PROJECT REPORT

Developing a Comprehensive, System-Wide Forecast to Support High-Penetration Solar

Gavin Newsom, Governor

September 2020 | CEC-500-2020-060

PREPARED BY:

Primary Authors:
Benjamin L. Norris, Marc Perez, Philip Gruenhagen, Jordan Hazari, Thomas E. Hoff (Clean Power Research, LLC), Frank Monforte (Itron, Inc.), Jan Kleissl, Elynn Wu, Rachel E. S. Clemesha (University of California, San Diego), Richard R. Perez, Jim Schlemmer, Sergey Kivalov (State University of New York, Albany)

Clean Power Research, LLC
1541 Third St., Napa, California 94559
(707) 258-2765
www.cleanpower.com

Contract Number: EPC-17-003

PREPARED FOR:
California Energy Commission

Liet Le
Project Manager

Jonah Steinbuck, Ph.D.
Office Manager
ENERGY GENERATION RESEARCH OFFICE

Laurie ten Hope
Deputy Director
ENERGY RESEARCH AND DEVELOPMENT DIVISION

Drew Bohan
Executive Director

DISCLAIMER

This report was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this report.

ACKNOWLEDGEMENTS

The authors acknowledge participation in the Technical Advisory Committee by Amir Javanbakht and Amber Motley (California ISO); Craig Smith, Alva Svoboda, and Mike Voss (PG&E); Nicholas Sette (SCE); and Brian D'Agostino (SDG&E).

The original GOES-15 data were obtained from NOAA, Comprehensive Large Array-Data Stewardship System (CLASS) (www.class.noaa.gov). The code used in the UCSD portion of the study is hosted at GitHub and can be accessed directly from https://github.com/elynnwu/Sc-line-forecast.

Clean Power Research, LLC. further acknowledges funding provided by the California Energy Commission, NOAA's Coastal and Ocean Climate Applications (COCA) program, DOI's Southwest Climate Science Center; San Diego Gas & Electric Company (SDGE) for access to their ground station irradiance measurements; helpful comments by Monica Zamora, Xiaohui Zhong, Handa Yang.

PREFACE

The California Energy Commission's (CEC) Energy Research and Development Division supports energy research and development programs to spur innovation in energy efficiency, renewable energy and advanced clean generation, energy-related environmental protection, energy transmission and distribution and transportation.

In 2012, the Electric Program Investment Charge (EPIC) was established by the California Public Utilities Commission to fund public investments in research to create and advance new energy solutions, foster regional innovation and bring ideas from the lab to the marketplace. The CEC and the state's three largest investor-owned utilities — Pacific Gas and Electric Company, San Diego Gas & Electric Company and Southern California Edison Company — were selected to administer the EPIC funds and advance novel technologies, tools, and strategies that provide benefits to their electric ratepayers.

The CEC is committed to ensuring public participation in its research and development programs that promote greater reliability, lower costs, and increase safety for the California electric ratepayer and include:

- Providing societal benefits.
- Reducing greenhouse gas emission in the electricity sector at the lowest possible cost.
- Supporting California's loading order to meet energy needs first with energy efficiency and demand response, next with renewable energy (distributed generation and utility scale), and finally with clean, conventional electricity supply.
- Supporting low-emission vehicles and transportation.
- Providing economic development.
- Using ratepayer funds efficiently.

Developing a Comprehensive, System-Wide Forecast to Support High-Penetration Solar is the final report for the Developing a Comprehensive, System-Wide Forecast to Support High-Penetration Solar project (EPC-17-003) conducted by Clean Power Research. The information from this project contributes to the Energy Research and Development Division's EPIC Program.

For more information about the Energy Research and Development Division, please visit the CEC's research website (www.energy.ca.gov/research/) or contact the CEC at 916-327-1551.

ABSTRACT

A broad range of new methods has improved the accuracy of solar photovoltaic fleet forecasting for California's electric grid. This improved accuracy ensures that new solar photovoltaic generation can be effectively integrated into the state's existing electric distribution and transmission grid. California ratepayers benefit from this improved accuracy by avoiding the costs of unnecessary fossil-fueled generation that would otherwise be required to make up for forecast uncertainties. This project developed methods for improved irradiance forecasts for California's coastal marine layer, including an innovative approach that models elevation-related dissipation. It quantifies improved regional- and site-specific blends of forecasts from alternate sources as part of its forecast horizon and addresses fleet variability by enhancing a previously developed method for probabilistic forecasting. To improve these planning tools, this project used as its foundation a publicly accessible database of production profiles of 504 actual solar photovoltaic systems located across California. The resulting distributed-energy-resource production database can be used for load studies that are scaled appropriately, based upon study requirements. New valuation metrics were developed and demonstrated to estimate savings from the project's improved forecasting methods. Finally, the project developed advanced regression methods that enable energy-resource planners to estimate future customer solar photovoltaic-generated capacity. These methods may be used by utilities to estimate future solar photovoltaic capacity, its technical load, and its financial impacts.

Keywords: Solar forecast, high-penetration solar, probabilistic forecasting, DER forecasting, forecast error, solar penetration, solar photovoltaic forecast.

Please use the following citation for this report:

Norris, Benjamin, et.al. 2020. *Developing a Comprehensive, System-Wide Forecast to Support High-Penetration Solar*. California Energy Commission. Publication Number: CEC-500-2020-060.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

Background

The California Legislature passed, and then-Governor Jerry Brown signed into law, Senate Bill (SB) 100 (De Leon, Chapter 312, Statutes of 2018). The 2018 landmark law mandates that renewable energy and zero-carbon resources supply 100 percent of electric retail sales to end-use customers by the end of 2045. It further requires the California Energy Commission, the California Public Utilities Commission, and the California Air Resources Board to prepare a report on the law's progress. Achievement of this ambitious goal depends upon speedy development of renewable resource generation and efficient ways to incorporate that generation into California's existing electric distribution and transmission grid. Renewable solar photovoltaic generation is a key component of that renewable energy mix.

Solar forecasting is essential to the multifaceted operation of California's electric grid. Forecasts of solar-generated energy are key to accurately scheduling supplemental resources so that supply and demand are balanced, and so demand is met in the most cost-effective manner. Forecasts are developed in the short term, up to a few hours ahead, and in the longer term: up to days, weeks, or even years.

Solar photovoltaic forecasting is performed for large, centralized facilities and small behind-the-meter resources on the periphery of the grid. While the methods themselves are similar for large and small solar photovoltaic resources, they are typically performed separately. This is because, from the operational point of view of the California Independent System Operator, large, central solar photovoltaic facilities are monitored directly while small, behind-the-meter generation is embedded within total load so that only net load is visible at the wholesale level.

Behind-the-meter solar photovoltaic systems present unique forecasting challenges. They represent a wide diversity of sizes, orientations, shading patterns, and equipment designs. They are located in all of California's diverse microclimates, including the coast with its episodic fog-like marine layer. The electricity generated by these many individual systems is not metered directly; generation from these behind-the-meter solar photovoltaic projects is instead a small, undifferentiated component of total load that does not pass through any meter, invisible to the utilities and grid operators. At the metering point, power flows to the customer and to the utility.

Given these challenges, present research focuses on behind-the-meter solar forecasting and how to more effectively integrate that power into net load forecasts.

Project Purpose

This project's goal was to improve upon existing solar forecasting methods and develop new ones with greater accuracy. Improved accuracy helps ensure that new solar photovoltaic capacity can be effectively brought online and seamlessly integrated into the state's electric grid. California ratepayers benefit by avoiding the cost of unneeded fossil-fueled generation that would otherwise be required in anticipation of forecast uncertainties. The methods developed in this research project will result in greater forecast accuracy, improved electric reliability, and grid integration of increasing amounts of solar photovoltaic capacity.

The California Independent System Operator must manage both forecast uncertainty and the related issue of solar photovoltaic energy production variability. Moving clouds can make solar energy production highly variable, especially in the fall when small, patchy clouds alternately shade and expose solar panels. The research and methods developed in this project address this variability by advancing a previously developed method that accounts for fleet variability through probabilistic forecasting.

Grid planning with high-solar photovoltaic penetration is further complicated by California's diverse microclimates and many variations in solar system designs, roof orientations, and shading patterns. Solar photovoltaic hourly output patterns depend upon these highly location-specific factors. This project helped address this issue by improving a publicly accessible database of production profiles of 504 actual solar photovoltaic systems located across the state. This distributed energy resource database can be accessed for load studies and appropriately scaled to study requirements.

This project also quantified the benefits of improved forecasting. Economic benefits to the California Independent System Operator — and, by extension, California ratepayers — depend upon the availability of metrics to evaluate forecast errors. New valuation metrics were developed and demonstrated that estimate savings from improved forecast methods.

Finally, future-year solar forecasting cannot be based solely on the existing fleet of solar photovoltaic resources; it also requires estimates for future capacity not yet installed in the state. This project therefore also developed advanced regression methods that enable planners to estimate future solar photovoltaic capacities. California's electric utilities can use these new methods to estimate future solar generation, its technical load, and its financial impact.

Project Approach

The project team consisted of Clean Power Research, University of California at San Diego, State University of New York, Itron, and consultant and former California Independent System Operator short-term forecast manager Jim Blatchford.

The University of California at San Diego developed a new forecasting method for tracking morning coastal fog. This method anticipates when and where fog dissipates in midmorning, exposing solar resources. The State University of New York advanced methods for blending alternative forecasts and developing region-specific forecasts. These methods proved to be more accurate than a single, general-purpose forecast. Itron developed methods for integrating behind-the-meter forecasts into net load forecasts. This research evaluated improvements to the neural networks used for net load forecasting. Jim Blatchford reviewed results, collaborated with the California Independent System Operator, and managed the technical advisory committee. Clean Power Research demonstrated methods of probabilistic forecasting in variable solar photovoltaic production, developed metrics for evaluating forecast errors, improved a database of behind-the-meter solar production, and developed methods for forecasting future customer adoption of behind-the-meter solar capacity.

The technical advisory committee included staff from the California Independent System Operator, Pacific Gas and Electric Company, Southern California Edison, and San Diego Gas &

Electric. The committee reviewed overall progress and provided the research team with valuable guidance in the project's developmental phase.

Project Results

The project team developed a broad spectrum of forecast-related improvements, described here.

University of California at San Diego Forecast Improvements

The University of California at San Diego developed a marine fog "line forecast" method that shows higher forecast accuracy in coastal Southern California than in either Central California or Northern California. Coastal topography likely plays a role in this discrepancy. For example, the simple topographic elevation distribution in San Diego favors the line forecast since it has more consistent weather across the forecast area. The lack of forecast accuracy at the immediate coast and the sharp gradient of cloud dissipation time within a few kilometers of the coast suggest that, at the immediate coast, local processes are important in determining when clouds dissipate. While the satellite cloud-motion forecast performs best at the immediate coast because of the frequency of days when clouds last all day, it is less accurate when clouds dissipate.

State University of New York Forecast Improvements

The State University of New York's new site-independent models (several related versions, designated v4x) shows systematic performance improvements at all locations and time horizons when compared with the current operational model (v4). About half of the performance gain is from a better model-blending fit with ground-location ensembles (v4 was fitted to a three-location subset). The other half of the performance gain came from adding predicted sunlight exposure levels for more accurate model blending. This improved site-independent model fully retains the user-ready applicability of the v4 version tested by the Electric Power Research Institute.

Location-specific forecasts exhibit additional performance improvements. Forecast skill exceeds 45 percent across all time horizons. For this site-specific model, much of the gain in performance is attributable to its use of site-specific parameters.

Probabilistic Forecasting

Clean Power Research's probabilistic forecasts were very promising because predicted and measured fleet output from a small solar photovoltaic fleet coordinated effectively with previously collected solar photovoltaic data. The probabilistic model accurately predicts fleet power variability in small geographic regions, even on days when fast-moving clouds cause high variability.

In some results there was a discrepancy between predicted and measured data. This is most likely because of the quality of available cloud-speed data. Further testing and data collection will be required to better understand this effect.

Distributed Energy Resource Production Database

Clean Power Research successfully demonstrated how physical system attributes can accurately infer measured solar photovoltaic output, developing inferred physical attribute

specifications that yielded excellent overall results. Clean Power Research modeled system outputs using both inferred system specifications and specifications reported by system installers. Median hourly error over a test year in simulations using inferred specifications was only 10.1 percent, while simulations from reported specifications had a median error of 16.7 percent.

Short-Term Load Forecasts

This study demonstrated improved accuracy for existing reconstituted load models by replacing unadjusted generation data with statistically adjusted data. Statistical adjustments allowed the measured load impact of solar photovoltaic generation to vary by load zone, time of day, temperature, and weekdays instead of by weekend days (that is, adjustments for the five weekdays are handled separately from the two weekend days). Alternative approaches were examined for implementing statistically adjusted solar photovoltaic generation values in reconstituted load models.

This study also provided evidence that weather-related impact on forecasting has evolved since 2013. To capture this change, time trends were introduced into reconstituted load model specifications that allow weather responses to evolve over model-estimation periods.

The day-ahead forecasts of morning- and evening-hour loads saw accuracy gains of approximately 10 to 20 percent in mean absolute percentage error, which is a measure of forecast accuracy. Midday and afternoon day-ahead load forecasts had modest improvements of 3 percent. Dawn hours had accuracy gains of 8 percent.

Valuation of Forecast Error

Three models of forecast error were developed and evaluated. Note that these three models are not directly comparable as they seek to quantify error in three different contexts.

- Resource adequacy. Solar forecast improvements developed in this project were quantified. If new gas generation were built expressly to comprise 90 percent of the 3-hour-ahead forecast errors across the 6.3-gigawatt behind-the-meter photovoltaic fleet in California, this metric indicates potential energy savings of up to $176 million over the next 30 years in terms of avoided new variability mitigation build-outs. Resource adequacy as defined in this work is a metric for evaluating the "long term" costs of building new generation assets to mitigate solar PV forecast errors.
- Real-time market correction. This metric is the most complex of the methods described in this report because it requires time-correlated market data. Because it requires this use of familiar data, however, it could ultimately prove more acceptable to grid operators and other utility stakeholders. Real-time market correction is a metric for value in the "short run," that is, in a time frame in which only the existing resources are employed.
- Perfect forecast. This metric penalizes the most extreme cumulative over-prediction events. It forces consideration of not only instantaneous power but also of energy across more extended periods. This is important because renewable energy is plentiful but doesn't share the "ramp-up on demand" characteristics of conventional forms of generation. This method also quantifies the incremental cost of storage if it is required to guarantee photovoltaic forecast accuracy. The perfect forecast is a long-run metric

assuming a scenario in which energy storage (not natural gas) is used to eliminate *all* forecast uncertainty.

Distributed Energy Resource Capacity Forecasting

Clean Power Research trained a regression model using historical customer data for 2014 and created a forecast for 2015. Predictor variables were cost-effectiveness, adoption density, and market segment. All three predictors were statistically significant. The forecast error was calculated as -1.8 percent, a strong confirmation that the model accurately predicts overall distributed energy resource adoption by residential customers.

Technology/Knowledge/Market Transfer Activities

Clean Power Research and its partners (Itron; the University of California at San Diego; Jim Blatchford; and the State University of New York) presented project results at several industry conferences, such as an American Meteorological Society meeting, the World Conference on Photovoltaic Energy Conversion and Institute of Electrical and Electronics Engineers Photovoltaic Specialist Conference, and Energy Systems Integration Group Meetings. Additional activities follow.

- One journal article has been published by the University of California at San Diego in the journal *Solar Energy*.
- Websites: The data set created under the distributed energy resource production database is available for public download on the California Distributed Generation Statistics website.
- The task report on *Short-Term Load Forecasting* was sent to several independent system operators and regional transmission organizations in the United States and in other countries. Other system operators and utility contacts presented project results.
- The publication of this report will provide additional project exposure to the industry.
- Clean Power Research and its partners held periodic technical advisory committee meetings with CAISO and the three IOUs. In addition, a separate meeting with CAISO early on in the project led to the inclusion of the real-time market correction metric in the project. The short-term load methods have been used by the load forecaster, and CAISO has reviewed project results.

Benefits to California

Lower Costs

Improved solar photovoltaic production forecasts will decrease net load forecast uncertainty. The cost of procuring both spinning and non-spinning reserves will also decrease. Clean Power Research is creating and implementing forecast-error-valuation methodologies that will allow users to estimate the costs of improvements in forecast accuracy. These methods are also being applied to forecast improvements developed in this project. This will lower energy costs by reducing the amount of stand-by generation required to mitigate forecast load uncertainty. Using the resource adequacy method, Clean Power Research estimates that the forecasting methods developed in this project will save California $176 million over the next 30 years.

Greater Reliability

The project team worked with the California Independent System Operator on an improved solar photovoltaic forecast that could be integrated into its operations with the twin goals of providing greater transparency for behind-the-meter solar photovoltaic systems and contributing to system reliability. The uncertainty of photovoltaic power generation adds to utility customer energy bills since additional reserves must be carried to accommodate this uncertainty. Based on the resource-adequacy method, which is designed to eliminate 90 percent of forecast uncertainty, the expected savings in the 3-hour forecast will be $28 per kilowatt of installed photovoltaics, or about 3 percent of the cost of utility-scale solar. This project quantified specific reductions in reserve requirements; for example, the required reserves reduce from about 95 Watts per kilowatt of installed solar photovoltaics to about 70 Watts per kilowatt in the 3-hour ahead market when using a site-specific forecast rather than the baseline forecast. Project results will also provide valuable knowledge to grid operators and California decision-makers who help the state plan for increased distributed energy resource grid integration while maintaining reliability.

Environmental Benefits

Project results will improve California's environmental footprint by decreasing fossil-fuel-based reserves required by the grid operator to accommodate solar photovoltaic forecast inaccuracy. By improving forecast accuracy, fewer resources will be required to meet net load uncertainty.

CHAPTER 1: Introduction

Overview

California's electric grid is integrating ever-greater amounts of solar photovoltaic (PV) generation into its distribution and transmission systems. This renewable generation includes both large, centralized, and small, behind-the-meter (BTM) solar resources.

Traditional generation and market design rely upon dispatchable resources that are technically able to generate electricity on demand. Market factors including the cost of individual generation, its location, and electricity demand determine the actual dispatch of these units. Scheduling these units depends upon these factors and load forecasts. California's dispatch process is managed by the California Independent System Operator (California ISO).

Solar PV generation changes this market design in fundamental ways. First, the California ISO must forecast net load, or the remaining load after metered, non-dispatchable renewable resources are accounted for. These are the large, central solar PV and wind facilities that generate wholesale electricity for utilities and are typically connected to the high-voltage transmission grid.

Secondly, the load changes with the integration of solar PV and other BTM resources (with rooftop solar PV being the most common in California). In order to forecast the final dispatch of conventional generation, the load must first be forecasted, followed by a BTM generation forecast. The difference is the net load delivered to electric utilities. Finally, wholesale renewable resource generation must be accounted for, which leaves the remaining load that must be served through the competitive marketplace of dispatchable units.

Forecasts are developed in the short term (up to a few hours ahead), as well as in the longer terms of days, weeks, or even years. While forecasting methods are similar for both large and small resources, they are typically performed separately. This is because BTM solar PV generation presents unique forecasting challenges.

As opposed to large, multi-megawatt central solar PV resources, BTM solar PV resources represent a broad diversity of sizes, orientations, shading patterns, and equipment designs. They are located in all of California's diverse microclimates, including the coast with its elusive fog-like marine layer. A portion of the energy generated by these resources is not metered directly, but rather contributes to load reduction without passing through a meter, technically invisible to both utilities and system operators. At the metering point, power flows in either direction.

Given these challenges, much present research focuses on BTM solar PV forecasting and how to most effectively integrate this generation into net load forecasts usable by the California ISO.

Project Purpose

This research set out to improve existing solar forecasting methods and develop new ones with greater accuracy. Improved accuracy helps ensure new solar PV capacity is effectively integrated into the state's existing electric grid. California ratepayers benefit by avoiding the cost of unneeded generation from fossil fuels that would otherwise be required to accommodate forecasting uncertainties. These methods together result in greater forecast accuracy, improved electric reliability, and greater integration of lowest-cost solar PV capacity. Most importantly, they also ultimately contribute to California's ambitious mandated timetable for reducing dependence upon fossil-fueled electricity generation.

As the state's grid operator, the California ISO must manage both forecast uncertainty and the related issue of solar energy variability. With transient clouds, solar PV generation can be highly variable, especially in the fall when small, patchy clouds alternately shade and expose solar PV panels. This project addresses that hurdle by enhancing a previously developed method to incorporate that variability into probabilistic forecasting.

Grid planning with high solar PV penetration is complicated by California's diverse microclimates and variations in solar PV system designs, roof orientations, and shading patterns. Solar PV resources generate very different hourly outputs depending upon these highly location-specific factors. One way the project addresses this issue is by improving a publicly accessible database of solar PV production profiles of 504 solar PV systems across the state. This distributed energy resource (DER) production database can be used for load studies and scaled to study requirements.

The project also quantified the benefits of improved forecasts. Economic benefits to the California ISO — and, by extension, California ratepayers — depend upon the availability of metrics to evaluate forecast error. The project team developed and demonstrated new valuation metrics that estimate savings from improved forecasting methods.

Finally, future solar forecasting cannot be based solely upon the existing fleet of solar PV resources; it also requires estimates of future solar PV capacity not yet installed. This project developed advanced regression methods that enable planners to estimate future growth of solar PV-generated capacity. These methods can be used by electric utilities to estimate future capacity, technical load, and financial impacts.

Project Approach

The present work covers a broad range of activities that support the unique requirements of net load forecasting.

- University of California at San Diego Forecast Improvements: The University of California at San Diego (UCSD) focused its efforts on the difficulties inherent in forecasting the behavior of stratocumulus clouds — the coastal marine layer — that do not follow the predictable patterns of inland clouds. Coastal fog may be present at sunrise and before "burning off" during the morning hours, exposing higher-elevation regions to sunlight earlier than lower-elevation regions. UCSD developed and tested an innovative line forecast model that accounts for this coastal phenomenon.
- The State University of New York Forecast Improvements: The State University of New York (SUNY) at Albany advanced several methods for the real-time net load forecast

system used by the California ISO. These included maximizing the blend of forecasts available from a variety of sources, including Clean Power Research's (CPR) own cloud-motion vector method for tracking cloud movements. In this project, SUNY investigated custom forecast blends based on region- and site-specific characteristics.

- Probabilistic Forecasting: CPR advanced existing methods for quantifying the variability of aggregate solar PV output across the state. While raw irradiance data is available from satellites on a 30-minute basis, fleet variability could now be determined at 5-minute intervals. This new method is based upon the relationship between cloud speed, the distance between resources, and time intervals.
- Distributed Energy Resource Production Database: Roof orientation and shading patterns at customer sites are extremely important for modeling and forecasting solar PV generation output. However, these are generally unknown at the individual system level. In this research project CPR perfected a new method for detecting design specifications and other modeling-related attributes using metered interval data like those in smart meter installations. These methods produced greater accuracy than specifications recorded by system installers. The method was put to use to fill in missing data from 504 solar PV systems across the state, and the resulting Distributed Energy Resource (DER) Production Database was made available to the public on the California Distributed Generation Statistics website.[1]
- Short-Term Load Forecasting: This study revisited a previously developed model called the Reconstituted Load Model, to correct systematic day-ahead forecast errors that arose partially from misspecification of the impacts of BTM solar PV generation. To determine specific model refinements, two initiatives were undertaken: the first by testing the assumption that 1 kilowatt-hour (kWh) of BTM solar PV generation reduces measured load by 1 kWh under all calendar and weather conditions, the second by determining whether the measured-load response to weather variations has evolved along with the growing penetration of BTM solar PV installations. A series of statistical models was used to estimate the weather-response function for the years between 2013 and 2015 and between 2016 and 2018. Several weather-response approaches were compared in this project. These approaches included adjustments to observed errors, where solar PV generation was added to net load to produce a reconstituted load, and another where the BTM solar PV forecast was fed directly into the net load network as an independent variable. Comparisons then determined whether overall weather response has evolved over the last five years.
- Valuation of Forecast Error: There is no industry-recognized method for quantifying the cost of solar forecast error. Such a method would, however, provide the basis for both determining the benefits of improved forecasts and evaluating alternative forecasts. CPR developed three methods for this purpose.
    1. Resource adequacy method, which reduces forecast error by calculating the cost of gas turbine capacity.

[1] "CSI 15-Minute Interval Data," https://www.californiadgstats.ca.gov/downloads/#_csi_15_id

2. Real-time market correction method, which calculates the cost of forecast error from historical pricing in California ISO markets.
3. Perfect forecast method, which estimates the cost of eliminating all forecast uncertainty through energy storage.

- DER Capacity Forecasting: This research project developed machine learning (a subset of artificial intelligence) methods that forecast future customer adoption of DER technologies, including solar PV. The model incorporates predictive variables such as payback period, proximity to other adopters, and customer market segment. These variables are indicators of adoption probability along with training on historical data sets, which further ensures proper model calibration.

CHAPTER 2:
University of California, San Diego Forecast Improvements

Introduction

With annual mean coverage of 22 percent for the ocean surface and 12 percent for the land surface, stratocumulus (Sc) clouds are the most common clouds on Earth. These clouds strongly reflect incoming solar radiation, and due to their low cloud height, they emit outgoing longwave radiation similar to that on Earth's surface. Sc clouds therefore have a strong net radiative-inhibiting effect on Earth's overall radiative balance. Sc clouds form in a shallow planetary boundary layer capped by a strong temperature inversion. The inversion limits the vertical mixing of warm dry air above and cool moist air beneath, which together prevents cloud evaporation. Geographically, the highest Sc cloud land coverage is in the mid-latitude coastal region next to eastern boundary currents where the temperature inversion is part of the warm dry descending branch of the Hadley cell (an atmospheric convection where air rises at the equator and sinks at medium latitudes).

Coastal California is an area of high Sc cloud coverage during the late spring and summer months when the semi-permanent North Pacific High is most intense. Sc clouds greatly influence the weather, water, and energy of Earth's ecosystem and have, for many years, been a topic of extensive research. In recent years, an aggressive renewable resource energy mandate in the State of California has attracted more than a half million BTM solar PV installations. As solar PV becomes an increasingly important source of generation for the electric grid, it is critical that utilities and system operators maintain reliable service while maximizing solar PV generation integration. With a majority of BTM solar PV sited along the densely-populated California coast, an accurate forecast of Sc clouds during the spring and summer months is important because those clouds substantially reduce solar irradiance.

Two types of methods are traditionally applied to solar irradiance forecasting, depending on forecasting horizons. Imagery-based cloud advection is used for short-term forecasting, such as ground-based sky-imager systems which determine intra-hour forecasting; while satellite cloud motion vectors (CMV) are used for forecasting up to five hours ahead. Traditional image-based cloud advection assumes that frozen clouds move in the direction of the regional wind field. While this assumption is generally true for a few hours, it loses validity for longer-term forecasts. For longer-term solar forecasting, ranging from hours-ahead to days-ahead, physics-based numerical weather prediction (NWP) is used. NWP uses current weather observations to solve a set of primitive equations that numerically integrate future weather. Forecast accuracy varies considerably depending upon time, location, and weather condition. Perez et al. (2010) found that hourly averaged satellite CMV forecast mean bias error (MBE) and root mean square error (RMSE), on an annual basis range from: 0.2 Watts per meter squared (W/m^2) and 104 W/m^2 in an arid region like Desert Rock, Nevada; to 30 W/m^2 and 159 W/m^2 in a semi-arid elevated place like Boulder, Colorado. Mathiesen and Kleissl (2011) found that NWP models generally under-predict cloudy conditions, causing over-prediction of solar irradiance. Recent studies have combined satellite images and NWP to better improve short-term solar

forecasting. For example, in addition to using traditional CMV techniques, researchers Arbizu-Barrena (Arbizu-Barrena et al., 2017) use a NWP to allow both advection and diffusion to the cloud index derived from Meteosat Second Generation.[2] This technique outperforms traditional CMV in areas with low topographic complexity but is less accurate in areas where cloud patterns are influenced by the terrain, as they are in coastal California.

NWP forecasts of Sc clouds in coastal California have been improved through better cloud initialization and by modifying inversion base heights in NWP. Imagery-based cloud advection forecasts have received less attention. Traditional satellite CMV forecasts do not accurately predict how Sc clouds move or dissipate over time, largely because Sc clouds do not typically follow widespread synoptic wind directions. An example of false Sc cloud advection by traditional satellite CMV forecast was recently reported in Miller et al., 2017. Sc clouds over land often form at night and reach maximum coverage before sunrise. During the day, Sc clouds dissipate because of solar heating at the surface (and the resulting surface sensible heat flux[3]), the cloud's solar heating, and entrainment of drier warmer air from above. When dissipation of Sc clouds is not considered, frozen cloud advection in satellite CMV often under-predicts solar irradiance.

The objective of this work was to improve solar irradiance forecasting during Sc cloud days, primarily through quantifying the clouds' dissipation times. A proposed Sc cloud-edge forecast (line forecast) using the Geostationary Operational Environmental Satellite (GOES) will improve solar irradiance forecasting by factoring in this cloud dissipation. The forecast tracks the inland edges of Sc clouds. The novelty of this method is that it tracks the evolution (dissipation) of stationary clouds, while standard cloud-motion approaches consider only the advection of fixed-shape clouds. This method combines physical insights into lower atmospheric cloud-top heights under a strong inversion with statistical methods. While applied here to Sc cloud forecasting in California, the research team expects cloud-edge tracking to be equally valid for other overcast stationary clouds such as coastal Sc cloud forecasts elsewhere, and inland fog forecasts. For example, fog and low stratus in Germany pose a challenge for transmission system operators. In addition to the low-stratus risk forecast system, designed for day-ahead warnings, line forecasts for short-term forecasting could also help inform decision makers.

Approach

Cloud Edge Line Forecast – Conceptual Motivation and Assumptions

In coastal California, the Sc cloud inland boundary edge elevation is typically at its maximum during the early morning. Clouds thicken and spread at night from longwave cooling but begin to thin when longwave radiative cooling is balanced by solar heating shortly after sunrise. In Southern California, the terrain rises nearly monotonically, (meaning it does not decrease) and

---

2 http://www.esa.int/Applications/Observing_the_Earth/Meteorological_missions/Meteosat/Meteosat_Second_Generation3

3 Sensible heat flux quantifies the rate of heat exchanged between the Earth's surface and the atmosphere and how that impacts the temperature that we sense (for example a hot surface tends to heat the air to increase air temperature (daytime), or hotter air tends to heat the surface to cool the air (nighttime).

peaks at about 1.5 kilometers (km) elevation 40 to 80 km inland. The eastern boundary of Sc clouds usually follows isolines of land elevation. Sc clouds extend inland up to where the land elevation reaches the inversion base height, and the inversion base height equals the cloud top height. Sc cloud dissipation occurs after sunrise as the air mass heats and becomes cloud-free with increasing elevation.

The line forecast correlates between land elevation and the Sc cloud's eastern boundary by extrapolating from the cloud edge elevation to predict the future cloud edge's location.

Cloud Dissipation Time

The GOES visible channel captures a new image every 15 minutes. At each satellite image time stamp, a visible reflectance cloud test was performed, and the eastern boundary of Sc clouds and land elevation determined. The median land elevation of all the points over the boundary represents the elevation at each time step. The time step of the line forecast model is also every 15 minutes. An example of the Sc cloud inland boundary moving towards lower land elevation is shown in Figure 1 at a one-hour time step.

Figure 1: Geostationary Operational Environmental Satellite Visible Images, June 14, 2016

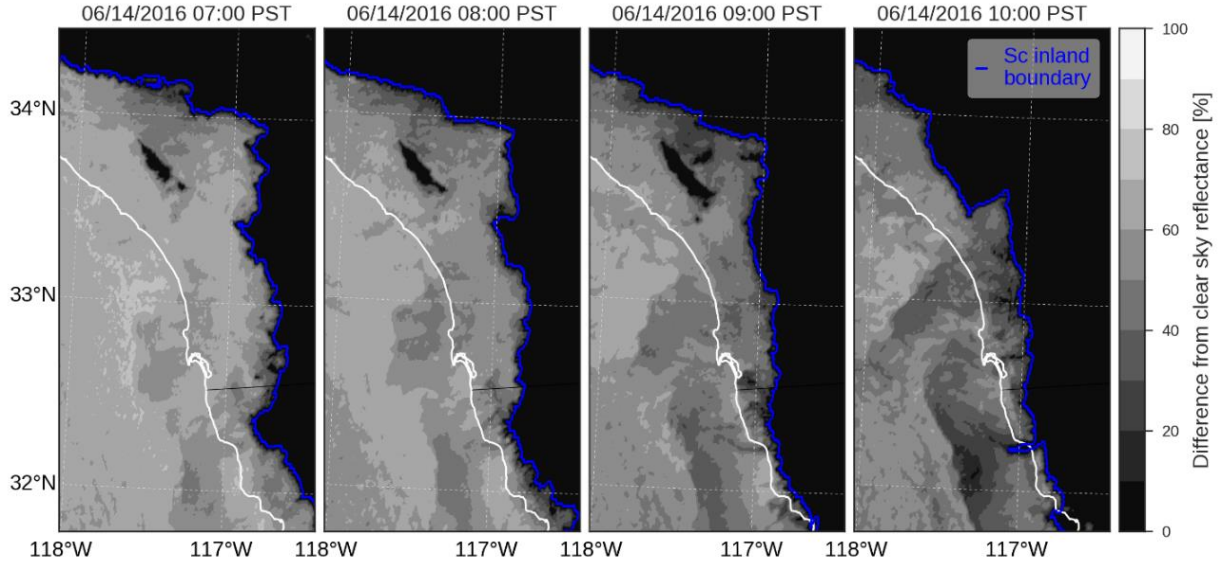

Sc cloud inland boundary highlighted in blue. Raw images are post-processed such that pixels within 15.5 percent of their clear sky reflectance are plotted as dark (clear), while pixels with larger than 15.5 percent difference are plotted in grey scales (cloudy). The position of the boundary advances towards the coast from 0700 to 1000 PST.

Source: Clean Power Research

Figure 2 shows the step-by-step approach used to issue a Sc cloud line forecast for this day. The accuracy of the line forecast is limited by the number of available images. While the forecast issued at 0800 Pacific Standard Time (PST) is accurate for this day, at 0700 PST the line forecast was unable to accurately predict the cloud dissipation; only two elevation points were available by 0700 PST, and the extrapolation of those median land elevations would have shown clouds persisting for the whole day.

Figure 2: Steps to Issue Line Forecast at 0800 Pacific Standard Time for Escondido

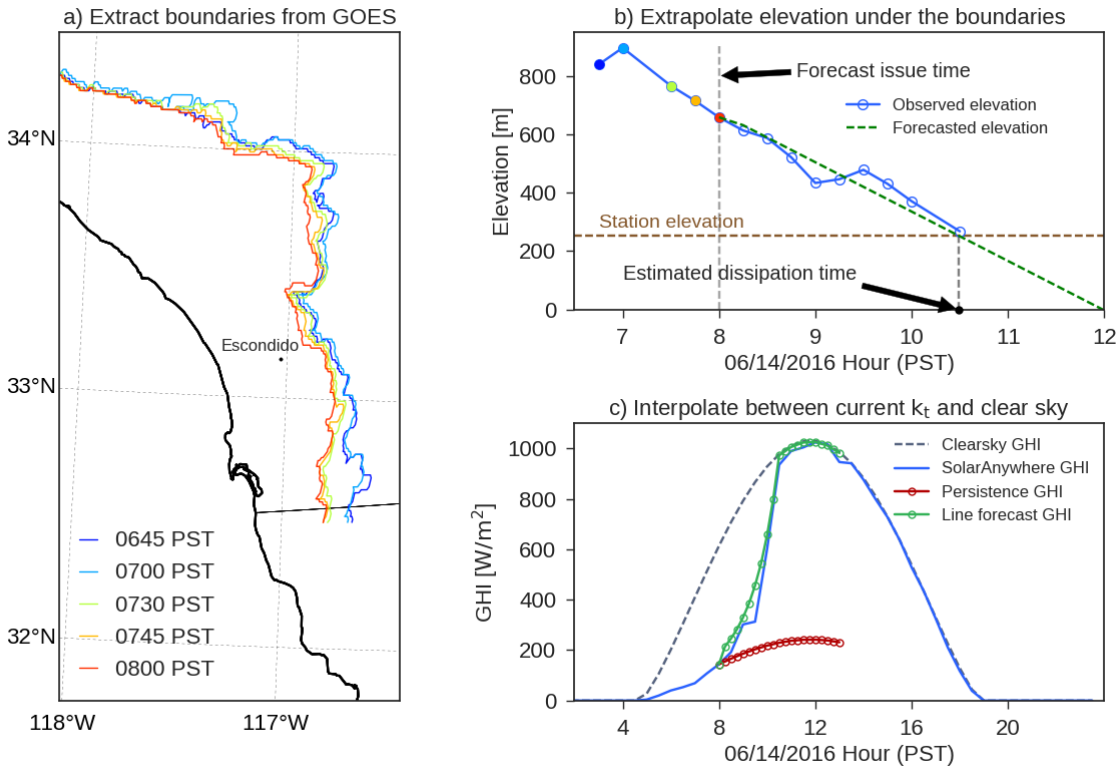

Steps include: (a) Extract Sc cloud inland boundary in consecutive GOES images up to 0800 PST and record the median land elevation under the colored lines (between 32.5ºN and 34.5ºN). (b) Extrapolate the median land elevation in time using a best fit line through 0645 to 0800 PST. Cloud dissipation time is when the forecasted elevation intercepts with the station elevation (here approximately at 1030 PST). (c) Interpolate between current k_t at 0800 PST and clear sky k_t at cloud dissipation time using the exponential function. The green line in (c) is the line forecast issued at 0800 PST with a 15-minute time resolution.

Source: Clean Power Research

To predict when Sc clouds will dissipate at a given location, a linear regression was performed on a time series of points corresponding to the Sc cloud eastern boundary land elevation, and extrapolated in time to the elevation for the given location.

Cloud Thickness Evolution

To describe the cloud thickness evolution between forecast issue time and dissipation time, the clear-sky index (k_t) is interpolated using an exponential function between forecast issue and dissipation times. When the sun angle is low (early morning), not much heat is received at the surface and the thickness of the clouds is approximately constant. As the day progresses, solar heating increases drastically, and the cloud thickness decreases more quickly. An exponential function was therefore chosen.

Results

Validation Against Satellite Observations

Southern California

Time series of hourly-averaged global horizontal irradiance (GHI) are shown in Figure 3 for forecasts issued from 0700 to 1000 PST. While the line forecast was computed for 19 days in June, this figure is the average of 5-15 June 2016 because of the availability of satellite CMV forecasts. Forecasts issued after 1100 PST are not of interest since clouds have already dissipated on most days and the line forecast would coincide with a persistence forecast with a clear sky index of 1. SolarAnywhere GHI shows that Sc clouds tend to stay longer at the coastal site where land elevation is lower. This is consistent with the assumption that the Sc cloud's eastern boundary moves toward the coast during the day as the average land elevation decreases.

Figure 3: Hourly Average of Forecasted and Satellite Observed Global Horizontal Irradiance

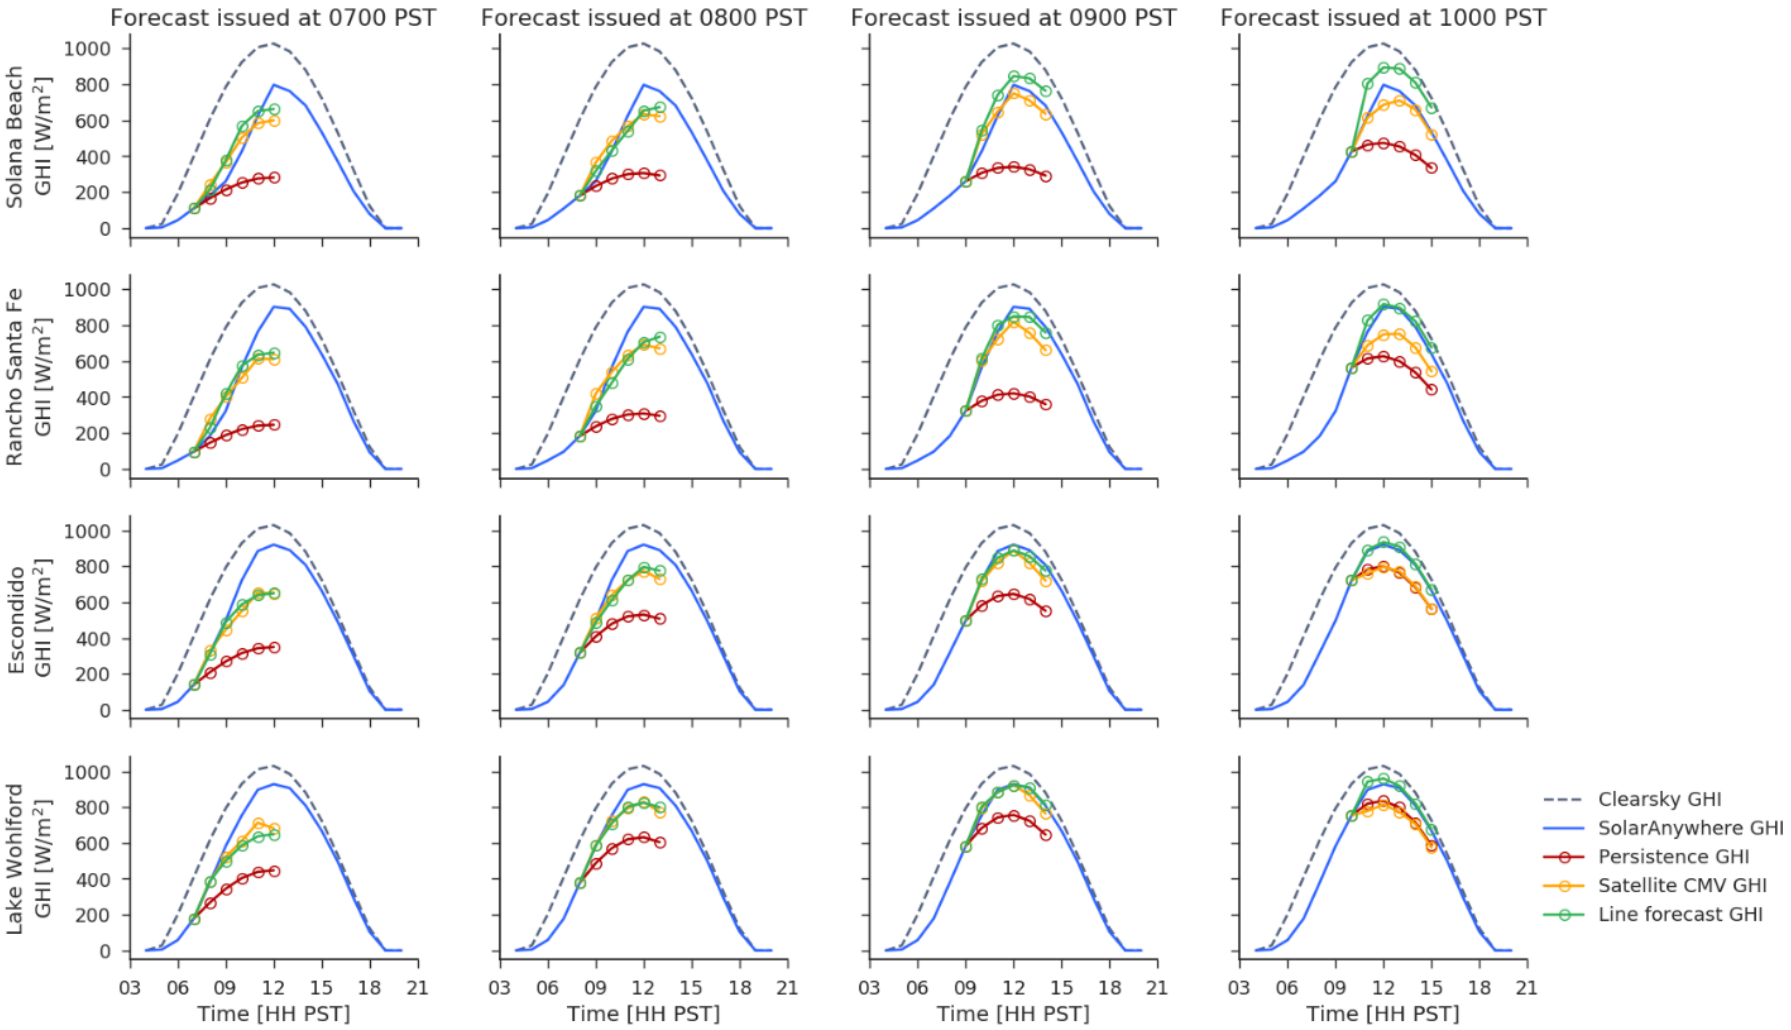

Hourly average of forecasted and satellite observed GHI for 5-15 June 2016 at Solana Beach, Rancho Santa Fe, Escondido, and Lake Wohlford. Each column represents a different forecast issue time. Note that each circle indicates the irradiance instantaneously at the hour, with the first circle corresponding to the real-time measured data, the second circle being the 1-hour ahead forecast, and the sixth circle being the 5-hours ahead forecast.

Source: Clean Power Research

MBE, MAE (mean absolute error), and FS (forecast skill) are shown in Figure 4. The line forecast consistently performed better than satellite CMV and smart persistence for all forecast horizon (1 to 5 hours ahead) and issue times at Rancho Santa Fe, Escondido, and Lake Wohlford. For Solana Beach, the line forecast was superior to the persistence forecast but slightly worse than satellite CMV. The line forecast had the lowest forecast skill for forecasts issued at 0700 PST. This is likely because only three visible GOES images were available at the time of the forecast issuance. In addition, dissipation often happens several hours after sunrise.

Figure 4: Error Metrics

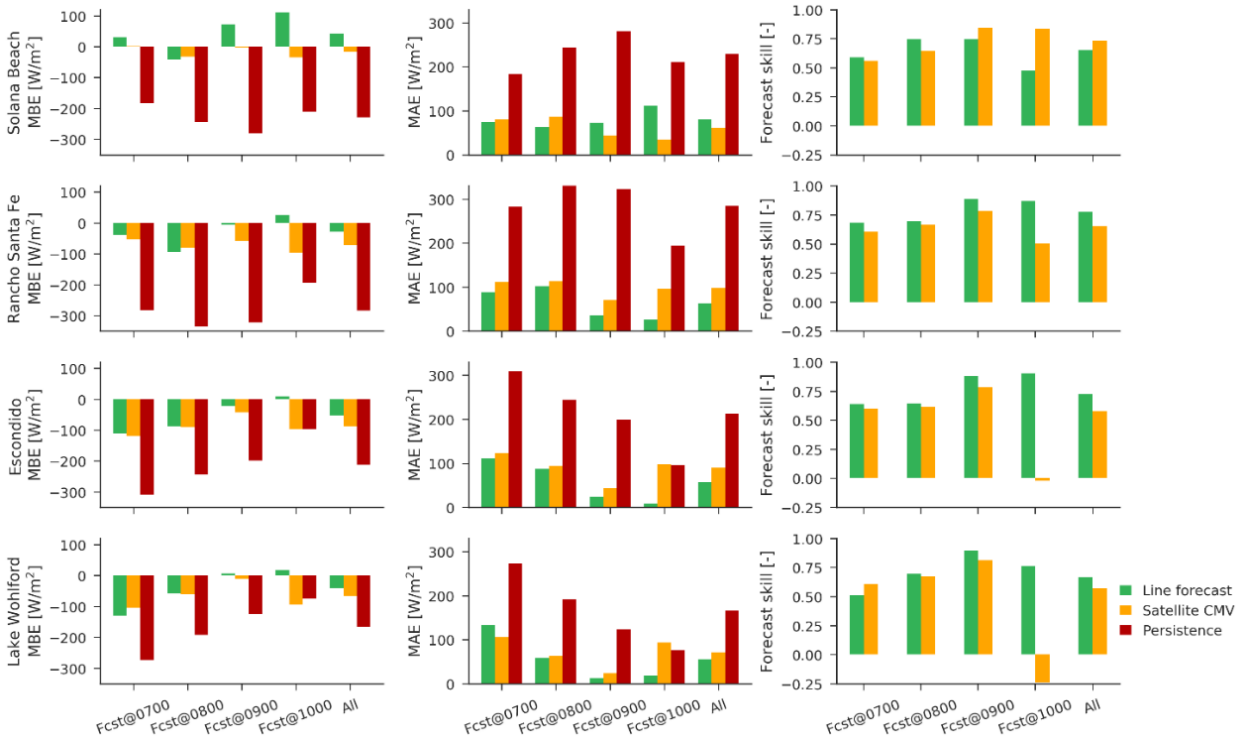

Averaged MBE, MAE, and FS for forecast horizons between 1 hour to 5 hours ahead for 5-15 June 2016 at Solana Beach, Rancho Santa Fe, Escondido, and Lake Wohlford.

Source: Clean Power Research

Geographical Error Distribution

To exhaustively quantify the usefulness of the line forecast, hourly SolarAnywhere GHI data at a horizontal resolution of 2 km were analyzed. Figure 5 is a spatial map of the line FS averaged over all forecasts issued at 0800 PST in Southern California, averaged across all forecast horizons, and Figure 6 shows all forecasts issued at 0800 PST in Central and Northern California. Note that FS has a maximum of 1, and that a positive value of FS represents an improvement over persistence forecasts. A time series of the boundary median elevation is shown in Figure 5.

Figure 5: Forecast Skill, Southern California

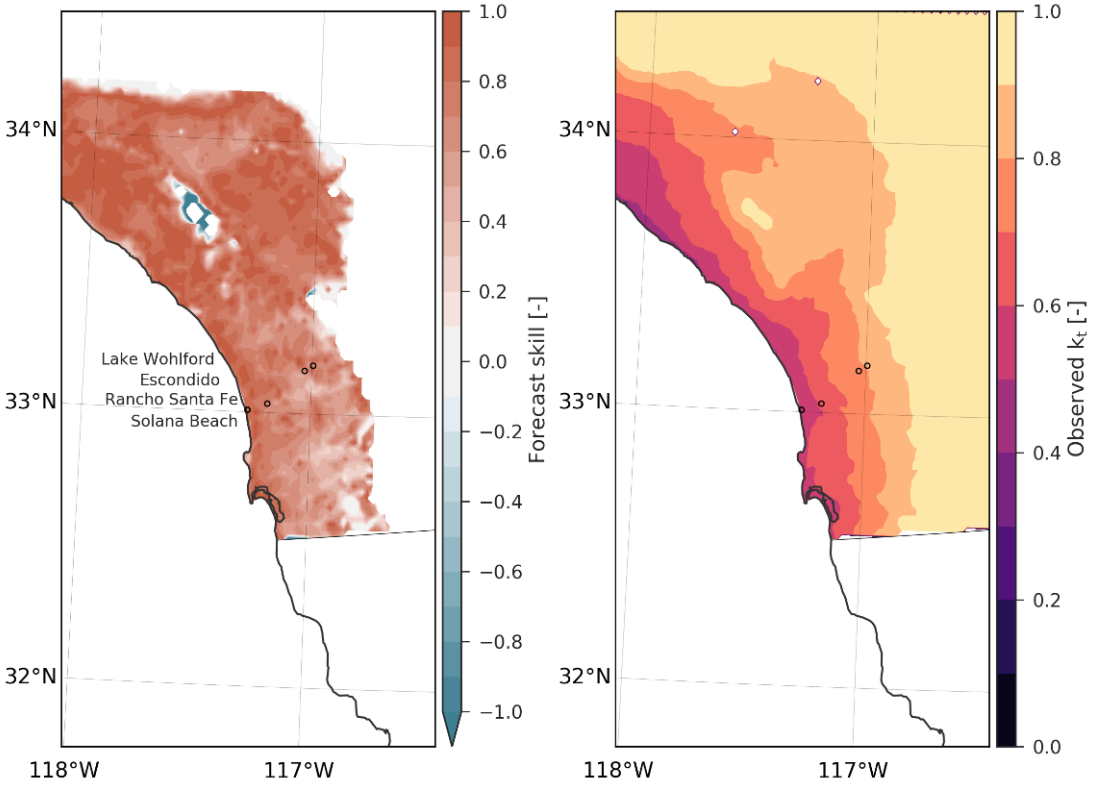

Left: Forecast skill of line forecasts issued at 8 am PST, averaged more than 1 hour to 5 hours ahead for 19 days in June 2016 for Southern California. Right: Satellite derived k_t averaged between 7 to 10 PST for the same 19 days.

Source: Clean Power Research

Figure 6: Forecast Skill, Northern California

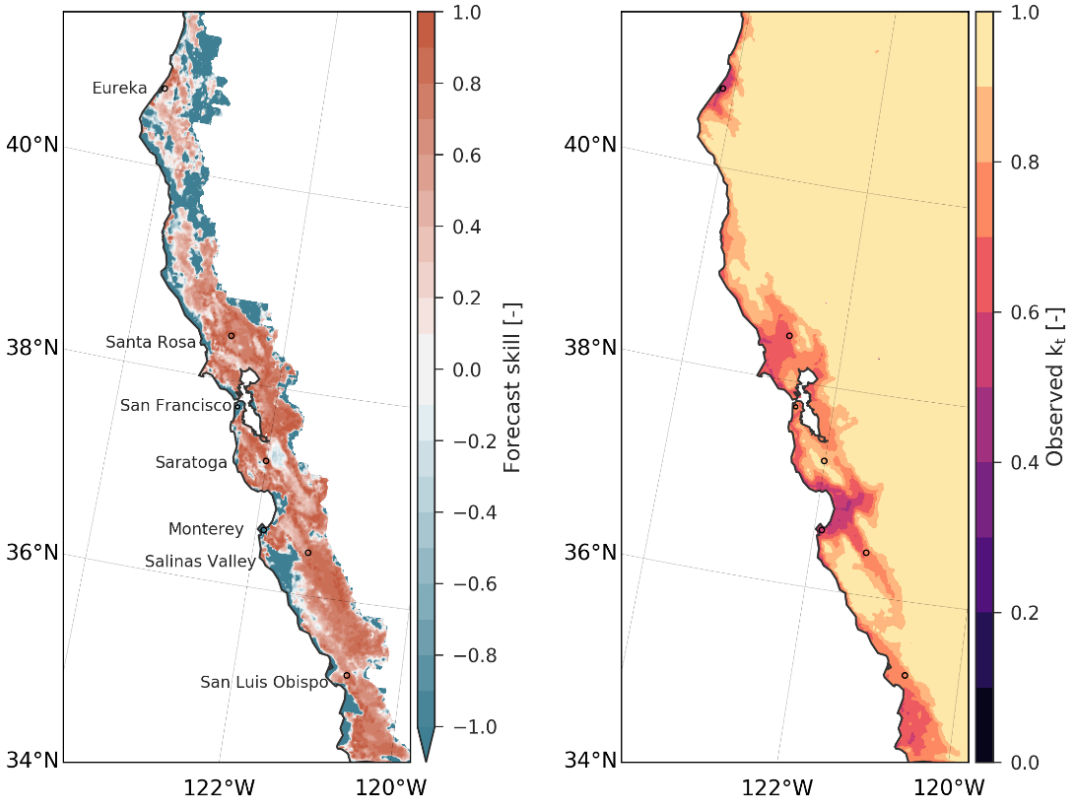

Forecast skill of line forecasts issued at 8 PST, averaged over 1 hour to 5 hours ahead for 25 days in August 2016 for Central and Northern California. Right: Satellite derived k_t averaged between 7 and 10 PST for 25 days in August 2016.

Source: Clean Power Research

Conclusions and Recommendations

For the two regions studied in this report, the line forecast shows higher forecast skills in Southern California than it does in Central and Northern California. The coastal topography likely plays an important role in this discrepancy. For example, the simple topographic elevation distribution in San Diego favors the line forecast since it has more consistent meteorological conditions across the forecast domain. The lack of forecast skill at the immediate coast and the sharp gradient of dissipation time within a few kilometers of the coast suggest that at the immediate coast local processes are important in determining when the clouds dissipate. While the satellite CMV forecast performed the best at the immediate coast because of the frequency of days when clouds persisted for the whole day, it was inaccurate on days when clouds did dissipate.

Improved Sc cloud forecasting is important because Sc cloud presence critically influences a broad array of solar applications. The novel Sc cloud forecast approach presented here may contribute to, among other applications, better management of utility electric grids and better planning for aviation. Future research will focus on greater understanding of factors that control whether Sc clouds persist during the day at the immediate coast.

CHAPTER 3:
State University of New York Forecast Improvements

Introduction

Pre-Sunrise Cloud Motion Vector

CPR's current CMV algorithm relies upon visible-band satellite images to both identify clouds and forecast their locations, based upon previous images. Unfortunately, there are no visible-band satellite images available before sunrise so forecasting using CMV is impossible in the early hours of the day. For this approach, CPR turned to infrared channels to identify clouds and forecast their formations.

Site- and Condition-Specific Tuning

The existing operational SUNY forecast model (Version 4, or v4) is based upon an optimized mix of models including a satellite-based CMV and global/regional solar- and cloud-cover NWP models.

Coupled with a PV simulation engine, v4 is readily used on any geographical scale, from single solar PV facilities to dispersed PV fleets, without requiring model training. It applies anywhere the underlying NWP and cloud motion models are available, basically 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 but quality training data are seldom available. However, it is important to stress that this key attribute does not rule out in-situ training when possible.

v4 was recently evaluated by the Electric Power Research Institute (EPRI) at test locations in North Carolina and California. It was the most accurate operational model among 13 US-based models from various providers. Several of the models were evaluated (for example, via deep machine learning) using measured solar PV plants data made available by EPRI.

CPR presented and evaluated a new model, SolarAnywhere Version (v4x), that retains the same philosophy (that is, readily deployable without system data training) but pushes operational accuracy further with better exploitation of operational model input data and utilization of SolarAnywhere historical data.

Approach

Site-Independent Enhancement

The existing operational v4 is derived from an optimized blend of a CMV model and four well-known NWP models: European Centre Mesoscale weather forecasts (ECMWF), Global Forecast System (GFS), National Digital Forecast Database (NDFD), and High Resolution Rapid Refresh (HRRR). The blend is a function of time horizon and time of day. The new version adds

predicted solar conditions as an additional input. The new optimum blend is determined empirically from 10 months of measured data at 7 Surface Radiation Budget Network (SURFRAD) locations. These locations sample a diversity of climatic environments in the contiguous United States.

Location-Specific Improvements

Historical satellite-derived irradiances are effective ground measurements to benchmark forecast models, produce comparable error metrics across climatically distinct locations, and forecast time horizons (Figure 7). CPR took advantage of this data to regionalize inputs and remove condition-specific biases. Historical SolarAnywhere data are available for long-time spans over entire continents and can be applied to locally optimize forecasts blends and biases in any region or location. This assumes, of course, that historical SolarAnywhere data are accurate enough and fully representative of a considered region or location. A growing body of evidence indicates that this is the case.

Figure 7: Root Mean Square Error Statistics

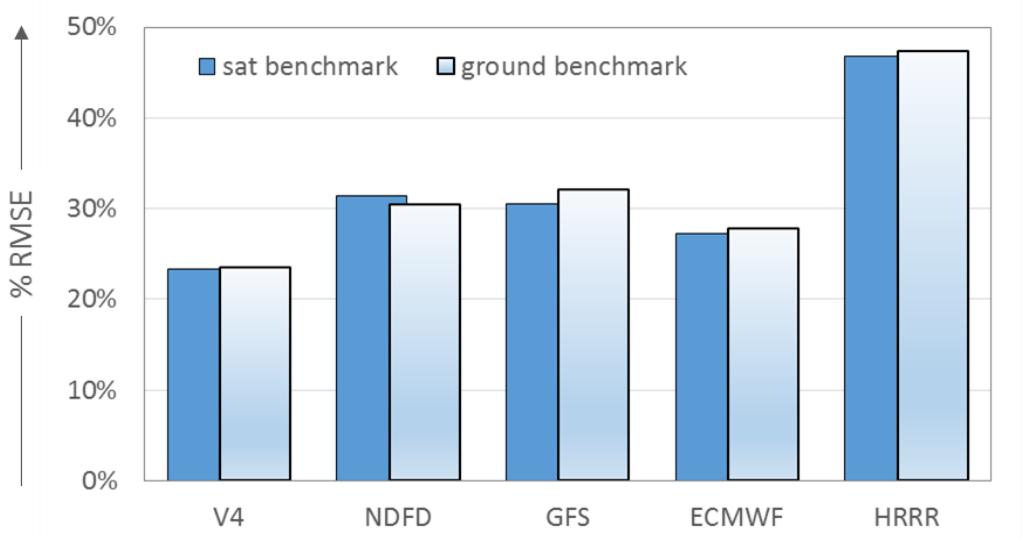

Comparing satellite and ground benchmarked forecast root mean square error statistics across all time horizons at multiple locations.

Source: Clean Power Research

Probabilistic Functionality

In addition to deterministic performance improvements, v4x also produces probabilistic information. The benefit of historical satellite irradiance data is the applicability for probabilistic forecasting without requiring ground-based measurements. Each deterministic forecast produces a probabilistic envelope from which operational probability quantiles can be readily deduced. The probabilistic envelopes are experience based, meaning that they are derived empirically from multi-site validations using either ground measurements or historical satellite data. CPR uses the same 10-month, seven-SURFRAD site sample as in the above blend to derive probabilistic envelopes.

Experimental Data

The data assembled for this study span from July 2015 to April 2016 and included:

- Hourly ground measurements for the SURFRAD sites of Bondville, Illinois; Boulder, Colorado; Desert Rock, Nevada; Fort Peck, Montana; Goodwin Creek, Minnesota; Penn State, Pennsylvania; and Sioux Falls, South Dakota.
- Historical intermediate-resolution SolarAnywhere Version 3 (SAV3) satellite-derived irradiances at arbitrary North America locations (hourly, 10-kilometer resolution).
- Satellite-derived CMV forecasts at arbitrary North America locations, with 1-5 hour time horizons.
- ECMWF, GFS, NDFD and HRRR NWP forecasts at arbitrary North America locations, with 1-48 hour time horizons except for HRRR (1-16 hours only).

Results

In Figure 8, CPR reports relative root mean square error (RMSE) 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 and smart persistence. The historical satellite model is also plotted as a visual baseline reference. Smart persistence follows the definition of the International Energy Agency (IEA) Task 36,[4] that is, accounting for both solar geometry detrending and increasing the integration time of the reference measurement commensurate with the forecast horizon. This definition is considerably more stringent (hence model taxing) than the commonly used, and often mislabeled, smart persistence that accounts only for solar geometry.

Figure 8: Relative Forecast Error

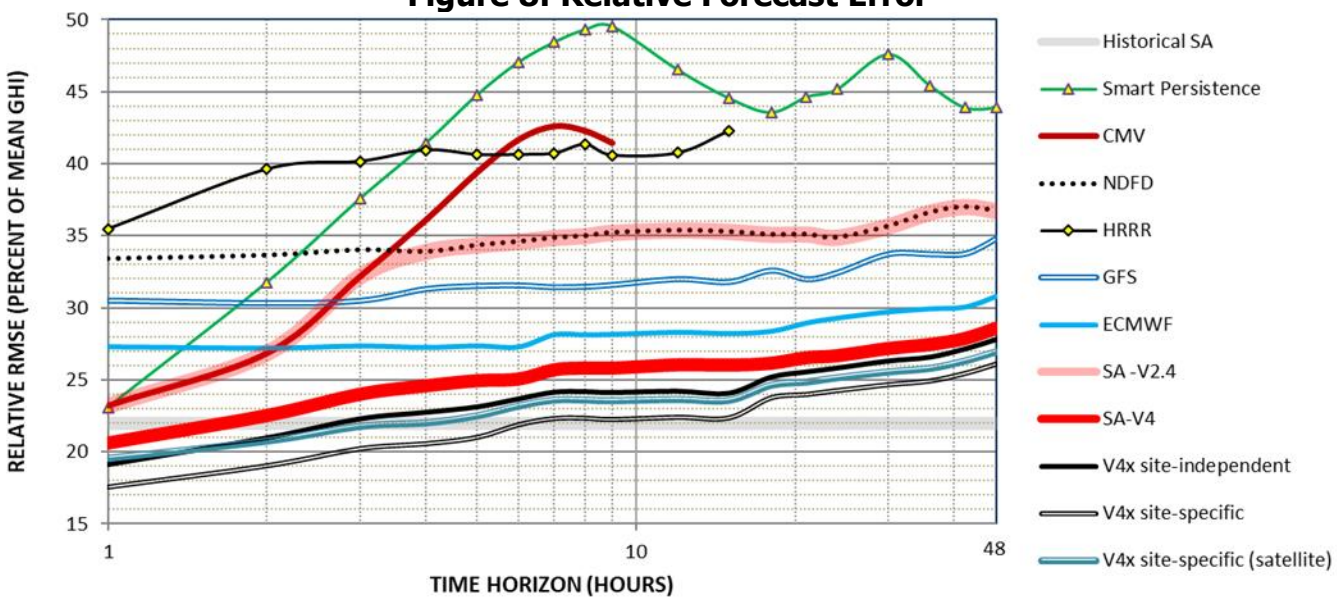

Relative forecast RMSE as a function of time horizon for the ensemble of seven SURFRAD locations.

Source: Clean Power Research

[4] https://iea-pvps.org/wp-content/uploads/2013/10/Photovoltaic_and_Solar_Forecasting_State_of_the_Art_REPORT_PVPS_T14_01_2013.pdf

Existing SUNY/SolarAnywhere forecast model versions include the original v2.4 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 9, CPR compares the skill of all versions at the operationally important 3-hour and 24-hour forecast time horizons. Skill is calculated in reference to IEA smart persistence.

Site-Independent Versus Location-Specific v4x Forecasts

The new site-independent model (v4x) shows systematic performance improvement at all locations and time horizons compared with the current operational model (v4). About half of the gain in performance originates from a better model blending fit with the ensemble of ground locations (v4 had been fitted to a three-location subset). The other half of 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.

Figure 9: Forecast Skill

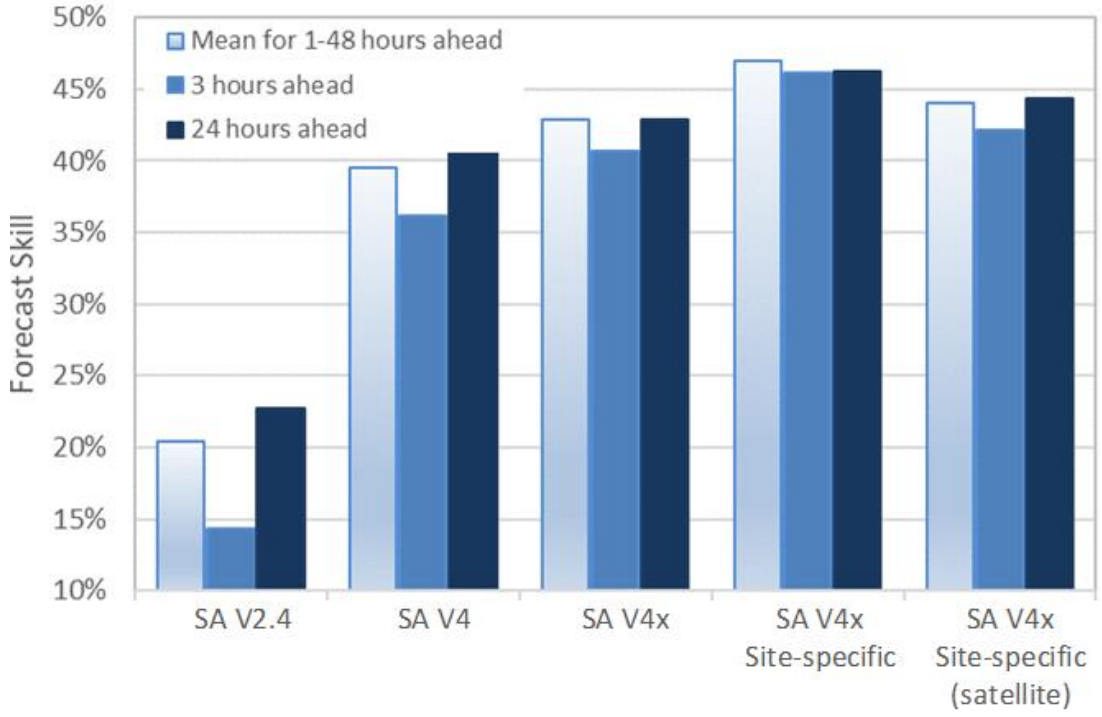

Forecast skill benchmarked to IEA smart persistence

Source: Clean Power Research

The location-specific version delivered further performance improvements. Forecast skill exceeded 45 percent across all time horizons. For this site-specific model, much of the gain in performance is attributable to its parameterization, which accounts for predicted insolation conditions. In Figure 10, CPR compares this performance improvement relative to v4: the light blue bars correspond to the site-independent model and the darker blue bars correspond to site-fitted models. The bars on the left represent model versions identical to v4 but fitted to all individual SURFRAD sites. The bars at right represent v4x models that optimize underlying model blends as a function of predicted insolation conditions. While insolation parameterization

increases performance steadily for the site-independent model, its 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.

Figure 10: Impact of Isolation Condition

Comparing the impact of insolation condition parameterization for site-independent and location-specific models.

Source: Clean Power Research

Operationally, the location-dependent forecast model performance represents a model that would be optimized locally from, for example, measured irradiances or PV plant output data. This localized performance improvement could be even further improved with real-time feedback techniques—for example, using machine learning or other techniques not addressed in this report.

Satellite-Derived Site-Specific Model

The key question here is how much of the site-specific performance enhancement discussed here can be captured without access to local measurements. SolarAnywhere historical data, however, is a validated proxy for local measurements.

CPR finds that using SolarAnywhere historical data to locally optimize model configuration captures some, but not all, of localized performance improvement potential. As shown in Figure 10, 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, which could bias model blending in favor of CMV for short durations. Nevertheless, systematic site-specific improvements were notable when achieved from the use of SolarAnywhere historical data (available locally throughout entire continents).

Application to regional forecast model optimization: CPR provided an example of site and regional model optimization for the State of California. CPR undertook this regional optimization work to develop operational forecasts for the climatic regions in California, shown in Figure 11.

Figure 11: California Energy Commission Climate Zones of California

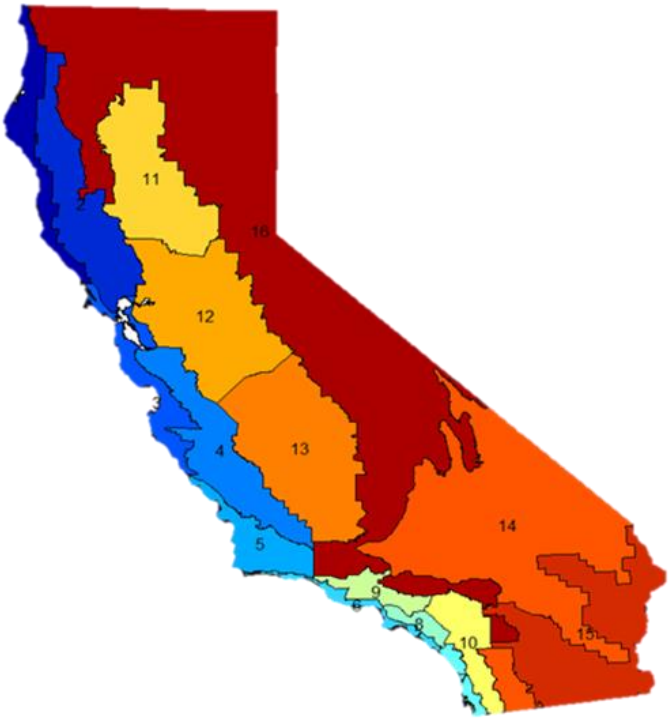

Source: Clean Power Research

This investigation augments the observations that (1) historical satellite irradiances are as effective as measurements for model validation, and (2) that satellite-based forecast training is an effective proxy for ground-based training.

CPR presents preliminary results for four climatically distinct California regions in this report.

- Region 1 - Northern Coastal Strip
- Region 9 - Los Angeles Hills
- Region 11 - Northern Sacramento Valley
- Region 15 - Imperial Valley Desert

Figure 12 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 with the site-specific results presented here. Models fitted from other California regions retain some of the performance gain, though when fitting based on other regions, there is little to no improvement and even in some cases some performance degradation.

Figure 12: Forecast Errors

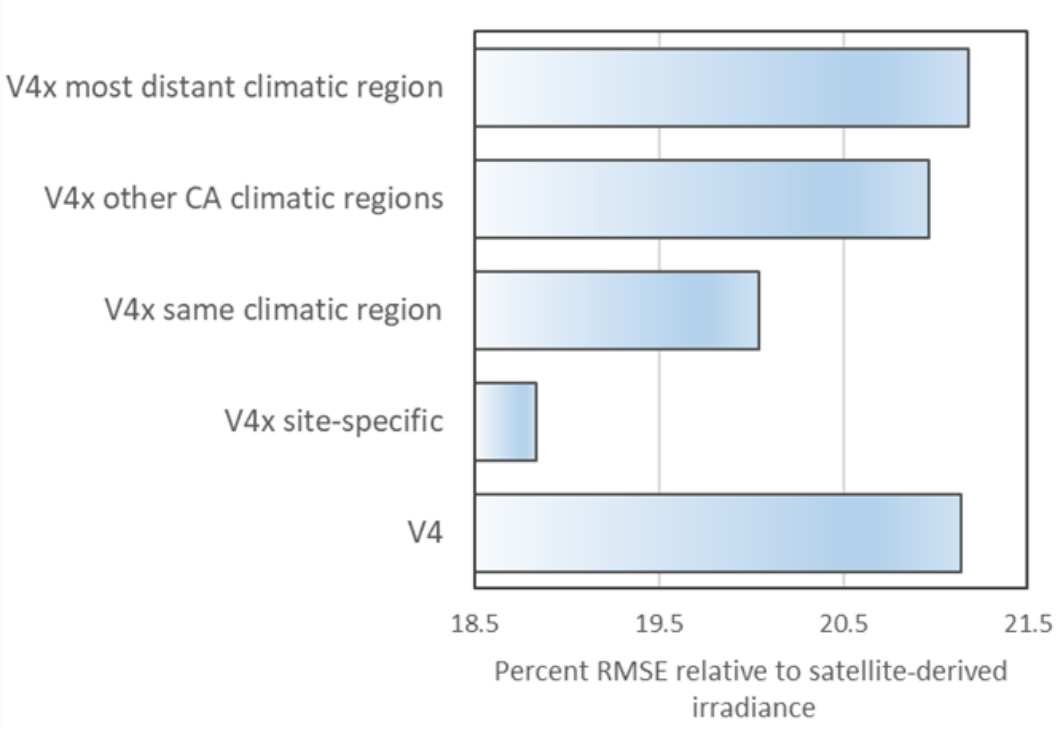

Resulting percent RMSEs for existing and locally fitted versions of the model across all forecast horizons (1-48 hours) for four California climatic regions

Source: Clean Power Research

Conclusions and Recommendations

CPR 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 showed measurable performance improvements over the current version.

Finally, CPR showed that model training using SolarAnywhere irradiances could capture substantial performance improvements with ground measurement training.

CPR recommends that future forecasts implement these model improvements. Implementation could include a locally trained model from either historical or ongoing ground measurements or from a locally trained model from satellite-derived measurements. Any of these options will both improve forecasts and reduce uncertainty.

CHAPTER 4:
Probabilistic Forecasting

Introduction

Electric grid operators increasingly rely upon solar PV generation forecasts when dispatching available energy resources. Solar PV forecasts include both large, utility-scale solar plants, and thousands of small BTM resources. In either case, the variability of solar PV energy output creates forecast uncertainties because of transient cloud cover. This is similar in some ways to the uncertainty of forecasting electric loads in general: the forecast can be considered a mean value, and its variability a measure of upper and lower bands of uncertainty.

Given these constraints for solar forecasters, this chapter describes a new way to quantify this variability. This method could, for example, deliver grid operators (such as the California ISO) a set of statistics to improve their operational decisions: the forecasted power output (mean output), the variance of power output (from which upper and lower bounds are derived), and the variance of power output (ramp rate). The ramp rate could apply to fast-acting ramping units.

Fleet variability is the variability from the aggregation of many systems. Fleet variability is low when systems operate independently, as when they are geographically apart with no output correlation, and operate as independent, uncorrelated systems. Fleet variability is more pronounced when systems are in close proximity, such as on the same distribution feeder. This characteristic suggests that load-serving entities could also use these probabilistic methods which could, for example, be used to better anticipate distribution-line fluctuations with increasingly high penetration of rooftop and community solar PV.

Approach

CPR provides SolarAnywhere® FleetView® to simulate and forecast average production for solar PV fleets. These simulations are typically based on satellite irradiance data that has a 1 km^2 spatial resolution and a 30-minute time resolution. FleetView estimates fleet output by simulating production for each distributed solar PV resource and aggregating its output. This means that PV systems within the same satellite data pixel (that is, the same 1 km^2 data point from the satellite image) use the same irradiance data to simulate energy production.

An important implication of FleetView's current design is that two identical PV systems (same rating, tilt, azimuth, shading, and so on) will have identical production profiles. This may be true when calculating the 30-minute average production for small, 1 km^2 pixels because short-term (such as 1-minute) PV power fluctuations are likely to average out over 30 minutes.

Simulating PV production using irradiance data that either has a much larger geographic resolution or a much smaller time resolution than 30 minutes may result in inaccuracies when performed for multiple resources sharing the same irradiance input data. Furthermore, producing variability statistics and probabilistic results may also be inaccurate since they require shorter-term data for variability calculations.

Inaccuracies are especially likely when simulating 1-minute solar PV production. FleetView produced identical 1-minute PV production profiles for two identical PV systems using the same irradiance data illustrated in Figure 13. Systems separated at a distance of less than 1 km, however, are unlikely to have the same 1-minute production under partly cloudy conditions, as illustrated in Figure 14. This type of data is used to calculate variability information.

Figure 13: Simulated Fleet Power from Non-Time Shifted Irradiance

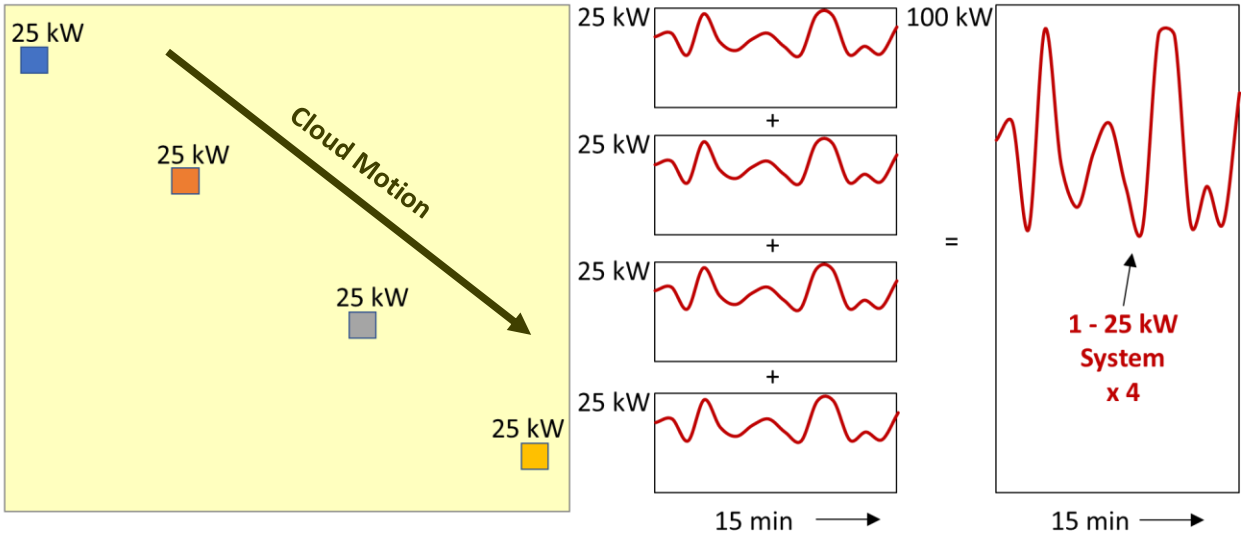

Hypothetical fleet power production of four systems, using identical irradiance data

Source: Clean Power Research

Figure 14: Simulated Fleet Power from Time Shifted Irradiance

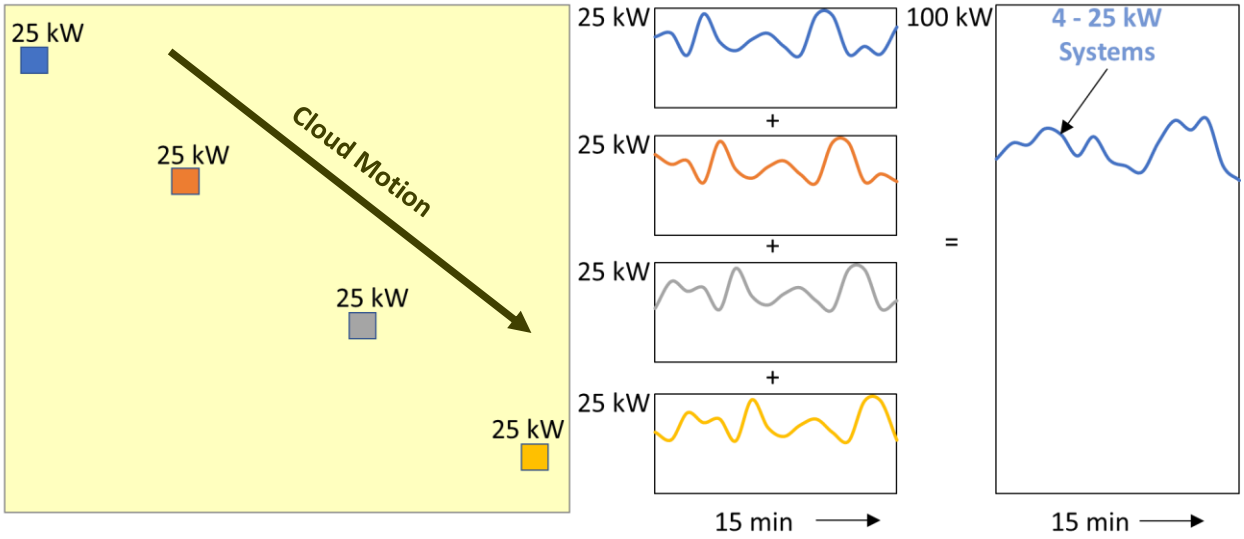

Hypothetical fleet power production of four systems, using individual irradiance data

Source: Clean Power Research

This issue is important because CPR produces a 1 km^2 spatial, 1-minute time resolution irradiance product. FleetView can use this product to simulate 1-minute PV energy production and from that simulation calculate the 1-minute variance of output over a 30-minute interval. There are, however, probable inaccuracies when simulating production using this data set. FleetView is likely to overestimate short-term variability in this instance for the reasons just described.

CPR has previously developed and patented a method designed to address this issue. This report describes a software implementation that employs the method and describes the processing steps involved when calculating probabilistic fleet forecast output. It also reviews the evaluation methods and presents their results.

While the method may be applied using spatial and temporal resolutions and averaging intervals other than those presented here, this project illustrates a method using SolarAnywhere FleetView capabilities. Spatial resolution is therefore approximately 1 km^2 (the resolution of the satellite imagery available to SolarAnywhere), and the temporal resolution is 1 minute (the temporal data resolution produced using CPR's cloud motion vector method). The averaging interval — the interval over which average power output and variance are calculated — of 15 minutes was selected for demonstration purposes. In other words, the power output "variance" represents the variance of 15 sampled data points.

Results

Dataset

CPR collected 10-second resolution irradiance data in small regions using 25 solar data loggers in November 2010 near Napa, California. These data loggers consisted of a small PV module attached to a device that measured and recorded voltage across a resistor. The voltage measurements were stored on USB drives and later converted to GHI.

CPR collected two sets of data. The first set placed the 25 data loggers within a 0.25 km^2 area. The second dispersed the 25 data loggers across a 25 km^2 area. The scaled and filtered data for each of the two data sets appears in Figure 15 and Figure 16. Both show 1-minute data for the 0.25 km^2 region and the 25 km^2 region after filtering the sites and scaling results using Clear Sky GHI from SolarAnywhere.

Figure 15: Time Series for Small Network

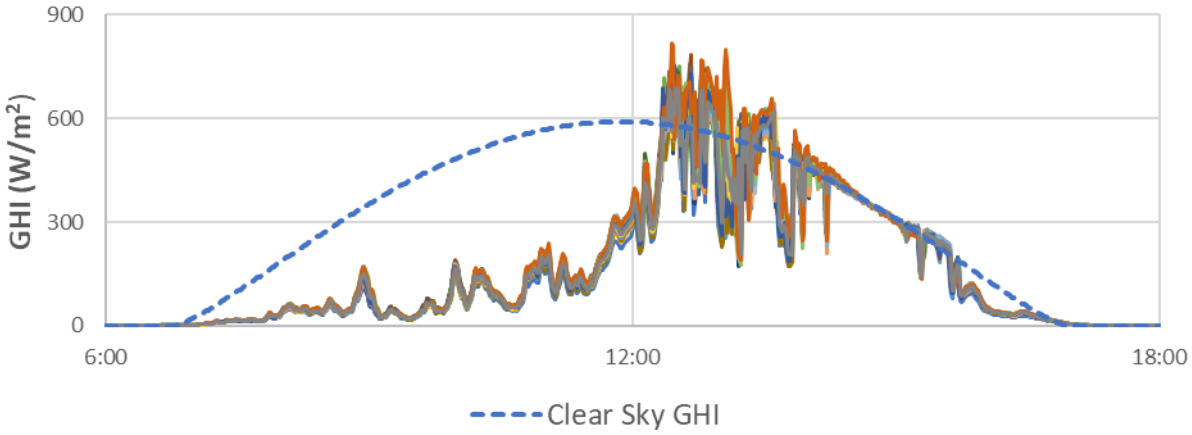

One-minute time series data for each site in 0.25 km2 network and Clear Sky GHI (11/7/2010)

Source: Clean Power Research

Figure 16: Time Series for Large Network

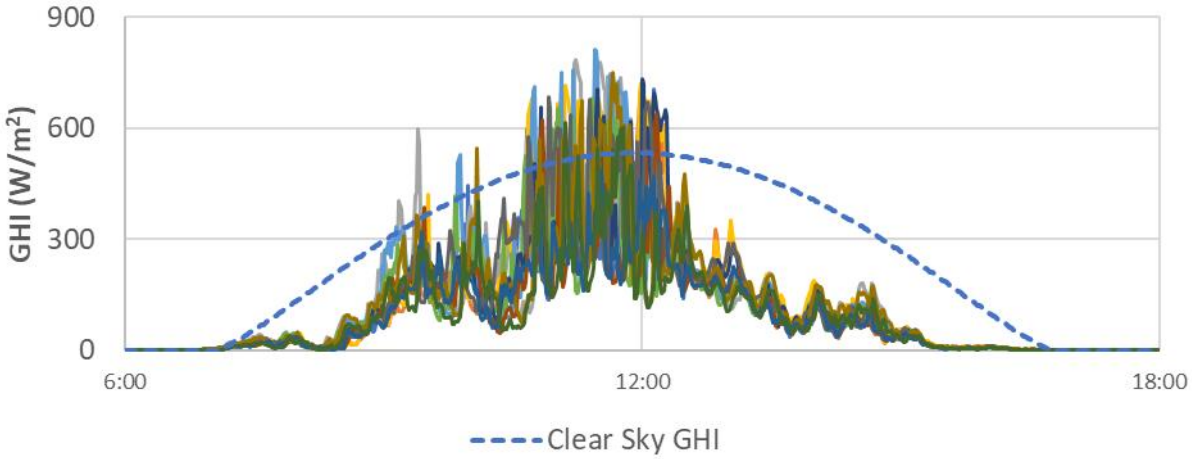

One-minute time series data for each site in 25 km2 network and Clear Sky GHI (11/20/2010)

Source: Clean Power Research

Figure 17 shows the key inputs required to calculate fleet power variability for the two networks. The left column corresponds to the 0.25 km^2 network and the right column corresponds to the 25 km^2 network. The first row is the Clear Sky Power, the second row the Kt Variance,[5] and the third row is the Correlation Matrix.

[5] Kt is the ratio of irradiance to clear-sky irradiance, that is, the irradiance that would be observed if clouds were not present. The values ranges from zero to one.

Figure 17: Required Inputs

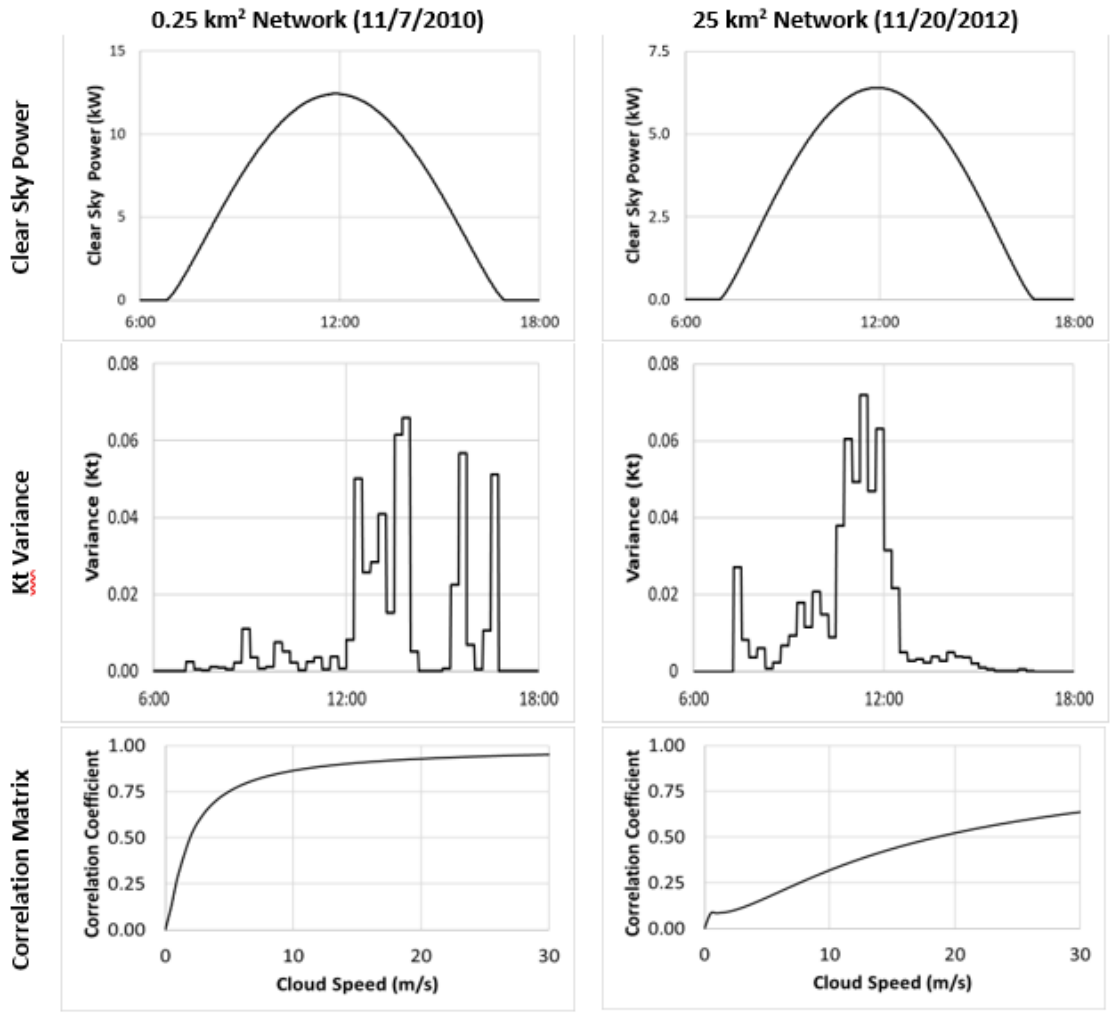

Inputs required to calculate power variability

Source: Clean Power Research

Evaluation

Evaluation results for each of the two datasets appear here. The days were selected based on high variability and minimal bad or otherwise invalid data. Three charts compare predicted output to measured output for each of the two datasets: Fleet Power, Standard Deviation of Fleet Power, and Standard Deviation of the Change in Fleet Power.

Figure 18, Figure 19 and Figure 20 correspond to results for the 0.25 km^2 region on November 7, 2010. Figure 21, Figure 22, and Figure 23 correspond to results for the 25 km^2 region on November 20, 2010.

Fleet Power

Figure 18 and Figure 21 compare measured fleet power with predicted mean fleet power. The figures include predicted upper and lower bounds (POE[5] and POE[95]),[6] and results indicate

[6] POE[N] is the probability of exceedance of the specified value N.

that measured fleet power almost always falls within the predicted upper and lower bounds. This is a good indicator that the probabilistic model is correctly predicting variability.

Standard Deviation of Fleet Power

Figure 19 and Figure 22 present standard deviation of fleet power throughout the day, broken into 15-minute time intervals. Results indicate that predicted standard deviation tracks very closely with measured standard deviation. This is another good indicator that the probabilistic model is correctly predicting variability. However, there are some periods where the predicted variability does not track as closely with measured variability. This is discussed in the Conclusions section of this chapter.

Standard Deviation of Change in Fleet Power

Figure 20 and Figure 23 present standard deviation of minute-to-minute changes in fleet power throughout the day, broken into 15-minute time intervals. Results indicate that predicted standard deviation tracks closely with measured standard deviation.

Figure 18: Small Network – Production

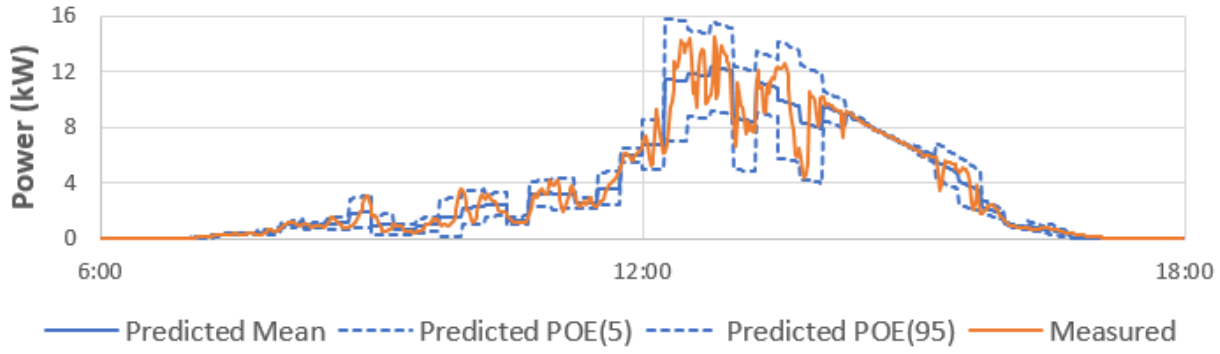

Production and variability statistics for 0.25 km^2 network of sites on variable day (11/7/2010)

Source: Clean Power Research

Figure 19: Small Network – Standard Deviation of Production

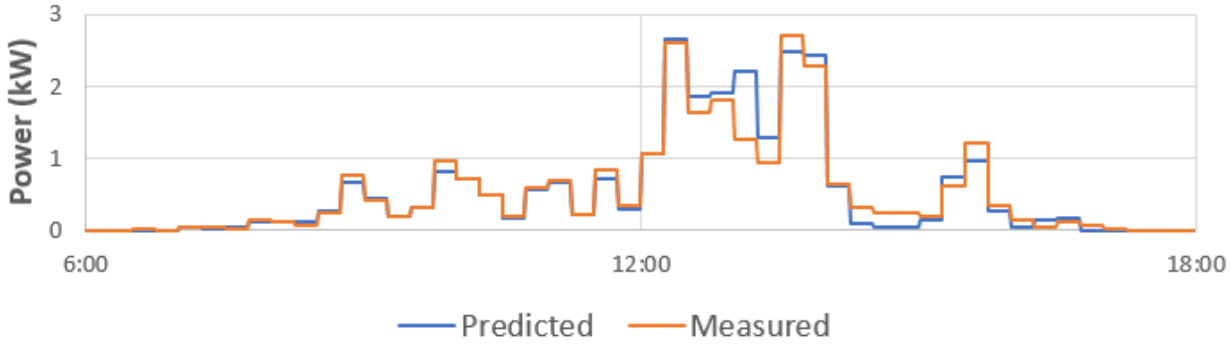

Standard deviation of power for 0.25 km^2 network of sites on variable day (11/7/2010)

Source: Clean Power Research

Figure 20: Small Network – Standard Deviation of Change in Power

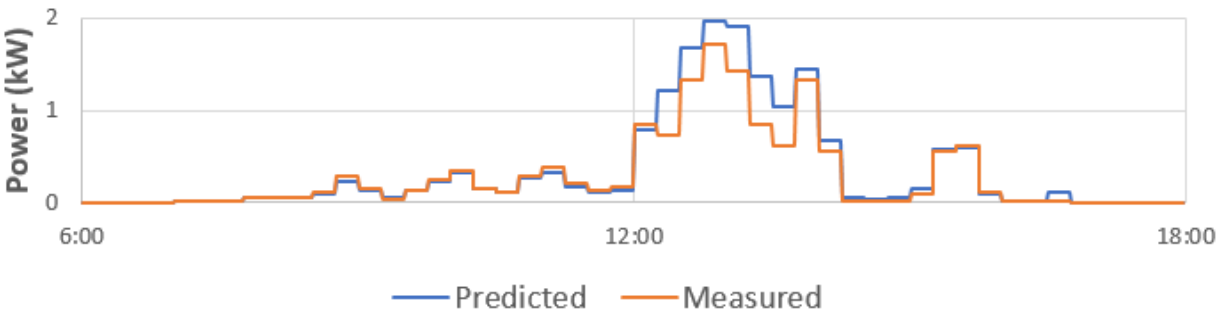

Standard deviation of change in power for 0.25 km^2 network of sites on variable day (11/7/2010)

Source: Clean Power Research

Figure 21: Large Network – Production

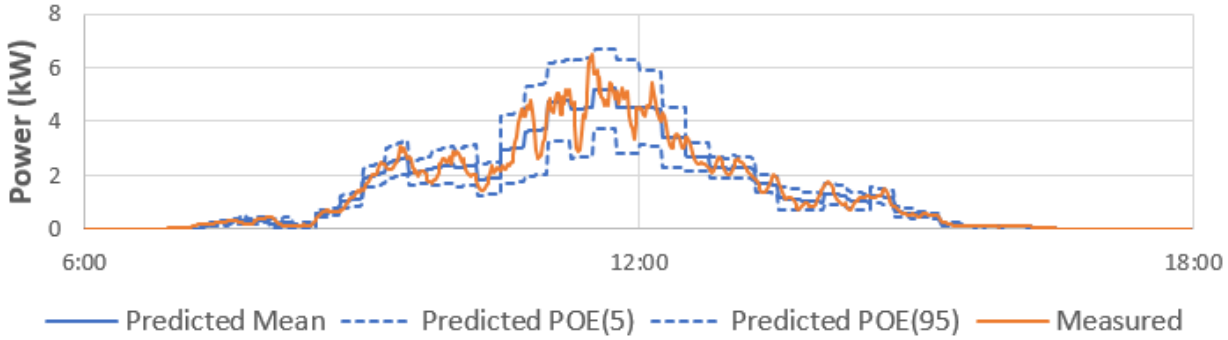

Production and variability statistics for 25 km^2 network of sites on variable day (11/20/2010)

Source: Clean Power Research

Figure 22: Large Network – Standard Deviation of Production

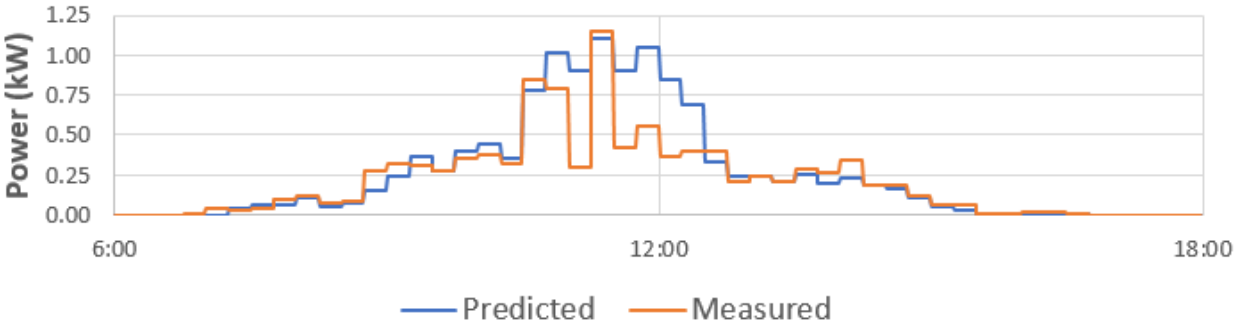

Standard deviation of power for 25 km^2 network of sites on variable day (11/20/2010)

Source: Clean Power Research

Figure 23: Large Network – Standard Deviation of Change in Power

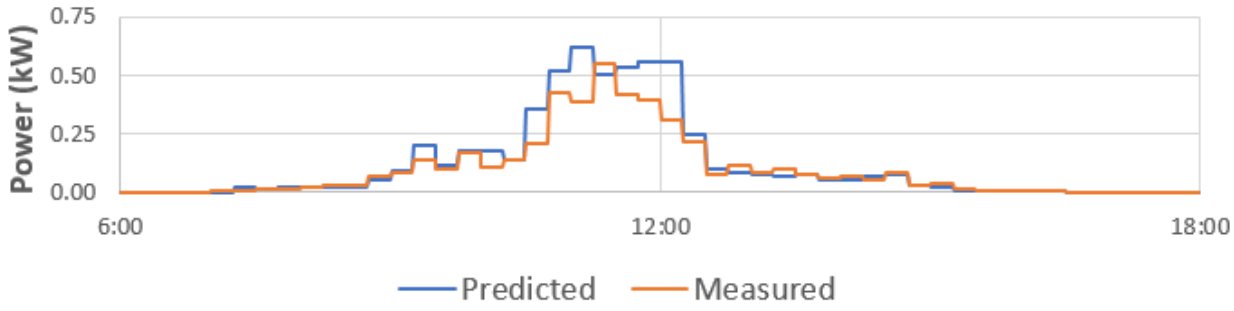

Standard deviation of change in power for 25 km^2 network of sites on variable day (11/20/2010)

Source: Clean Power Research

Conclusions and Recommendations

The results of these evaluation tests are very promising. The probabilistic model accurately predicted fleet power variability in small geographic regions even on days where fast-moving clouds created high variability.

There were, however, some periods where predicted results did not match measured results as closely as others. In particular, those periods occurred during the middle of the day for the 25 km^2 site network. It is difficult to determine the cause of this discrepancy due to the small data sample size. The most likely cause was the quality of cloud speed data available in November 2010. This is something that could be narrowed down and potentially improved with further data evaluation. It is also possible that the probabilistic model is over-predicting variability, especially when dealing with sites that are further apart, as in the 25 km^2 network of sites.

Next steps include:

- Software Enhancements. The most important task accomplished in this project was to convert equations in a patent-to-prototype software code. More work could be done improve the software implementation. The following tasks include potential work that could improve software performance.
- Improve SolarAnywhere Response Time. A few changes could be made to SolarAnywhere that would make response times faster, specifically for the types of data requests required for the probabilistic fleet forecasting model. In the general-use case for this model, a fleet would consist of many PV systems located closely together. SolarAnywhere irradiance data is represented at a 1 km^2 spatial resolution, which means the research team could potentially make optimizations by retrieving irradiance data just once per 1 km^2 pixel (rather than once per PV system). This would improve the general performance of the probabilistic model by speeding up the time it takes to retrieve irradiance data.
- Improve Overall Efficiency. In its current state, the software is not as efficient as it could be when performing some processing steps. For example, some of the fleet configuration information is calculated from scratch every time the program runs. Software could be designed to just look up the necessary information instead of devoting resources to calculate it.

- Use a Machine Learning Approach to Obtain Irradiance Statistics. Currently, the probabilistic forecasting algorithm uses input variability statistics based on 1-minute SolarAnywhere irradiance data. It is feasible, however, to apply a machine learning approach to generate 1-minute variability statistics based on 30-minute irradiance data and other variables (cloud speed, temperature, time of year, and so on). This way, the algorithm would not be limited by the restriction to generate 1-minute irradiance data.
- Integrate the Probabilistic Model into SolarAnywhere. Make the option to generate probabilistic fleet output available through the SolarAnywhere application program interface. This way, running the probabilistic model for fleet output will be both more convenient for users and standardized with other SolarAnywhere application program interface methods.
- Further Evaluation. A clear benefit of the probabilistic model is predicting power variability for a fleet with solar PV systems that are close together on high-variability days. An ideal evaluation dataset would include 50-100 residential PV systems located on a distribution feeder within 1 mile of each other.
- Examine Alternative Data Sources. As stated previously, the Napa, California, dataset was collected in 2010. The present evaluation therefore relied upon cloud-speed data generated from a legacy SolarAnywhere cloud motion code from 2010. A next step would be to generate cloud-speed data using the up-to-date SolarAnywhere cloud motion code for this time range. Cloud speed plays an important role in the probabilistic fleet output so it would be interesting to see how the results may or may not change using more accurate cloud-speed data.

CHAPTER 5:
Distributed Energy Resource Production Database

Introduction

CPR performed a broad study of solar PV forecasting methods while both developing a library of DER production data and refining DER growth projections in distributed resource plans (DRPs) developed in 2015 by California's three investor-owned utilities (IOUs).

The DER production database covers the target period from January 1, 2011, through December 31, 2016, and contains production data at 15-minute intervals for each zip code. The library is built on measured historical production data from 414 of 504 systems monitored under the California Solar Initiative (CSI) program. The remaining 90 of 504 systems were discarded because of technical data issues. CPR collected, processed, and analyzed the measured data, analyzed the error in modeled versus measured data, improved the measured data by combining it with modeled production data, then aggregated the data by zip code.

CPR also used the solar PV net energy metering Currently Interconnected Data Set (CIDS) to evaluate historic DER capacity growth and compared it with capacity growth estimations in the DRP prepared by each of California's three IOUs: PG&E, SCE, and SDG&E.

Approach

Overview

To create the database for DER production, CPR first collected, analyzed, and processed existing PV production data from CSI. As expected from a collection of this type, CPR found large amounts of missing or erroneous data. Therefore, to provide a complete data set for the target period CPR had to fill in the missing data with simulated (modeled) data. CPR performed simulations using both its SolarAnywhere® database of historical irradiance and temperatures from the locations and time periods considered.

CPR obtained design specifications for each system (system attributes required for modeling output) from two alternate sources: specifications reported by installers through the CSI program, and specifications inferred from periods of measured data. This inference of system specifications was one of the major efforts of this task.

CPR used these specifications (from either source, depending upon relative quality) to model PV production, then compared the accuracy of that modeled production data with the measured data. Finally, CPR combined the measured and simulated production data and aggregated it by zip code to create the final database.

Data Collection, Preprocessing, and Analysis

In 2010, CSI authorized its program contractor, Itron, Inc., to install production meters on 504 CSI systems. CPR acquired 15-minute-interval data for these systems from the California Distributed Generation Statistics web site. This measured PV production data formed the foundation of the DER production database.

Inferring Specifications From Measured Production Data

CPR developed a method to infer system specifications using measured production data prior to this project. This method determines the specifications which, when used to simulate production, resulted in lower hourly error levels when compared with the reference measured time series. The method was developed using a spreadsheet format for a single system. For this current research project, a one-by-one use of the spreadsheet was impractical for 504 systems, so CPR encoded the method in custom software for batch processing.

CPR's method employs the golden-section search (a technique for finding the minimum or maximum of a function inside a specified interval) to identify, one PV system at a time, values that minimize error. The method works through the attributes in a specific order by initially selecting a coarse approximation before refining some of its values in a second iteration. For example, in the initial modeling, CPR used a "constant horizon" (solar obstructions at the same elevation angle for every azimuth) to approximate both diffuse and direct losses from obstructions. Elevation angles for azimuth-specific obstructions were not selected until the values for other attributes had already been determined.

This method also stores results of previous simulations to avoid modeling the same specifications multiple times. Using this approach, CPR modeled a maximum of 431 different candidate specifications for each system for the period covered by the measured data and inferred specifications for 414 of the 504 systems. Using a combination of automated filters and manual inspection, CPR eliminated 90 systems because of bad or missing data.

Results

Modeling Photovoltaic Production

Two alternative sets of system specifications were available to model missing periods: specifications reported by installers, and specifications inferred using the method just described. CPR modeled every system over the target period of January 1, 2011 through December 31, 2016, using both sets of specifications.

SolarAnywhere simulations use a power model that calculates both direct and diffuse irradiance as it strikes the tilted plane of a PV system's modules, taking into account any solar obstructions, module efficiency changes from temperature and other DC losses, and the inverter's maximum power rating and varying efficiency at different power levels. The model can simulate both fixed, single-axis tracking, and dual-axis tracking systems, as well as systems with multiple arrays with different orientations and DC-to-AC ratios. However, the code that inferred specifications from measured data did not attempt to determine whether a system had multiple arrays with different orientations. Therefore, the inferred specifications were always single-array systems, with output that best approximated the measured data.

Comparing Results From Inferred and Reported System Specifications

After filtering the measured data, CPR calculated the hourly relative mean absolute error (rMAE) of the measured production data compared with data modeled using reported specifications and data modeled using inferred specifications. The rMAE for data from inferred specifications ranged from 6.3% to 50.8% with a median rMAE of 10.1%, while the rMAE for data from reported specifications ranged from 7.2% to 127.6% with a median rMAE of 16.7%.

For 96 percent of the time where the research team had both inferred and reported specifications, the modeled energy production data using inferred specifications resulted in a lower error rate (when compared with measured data) than the production data modeled using reported specs. Only 17 systems had lower error using reported instead of inferred specifications.

Figure 24 compares the distribution of error (as measured using rMAE) for production data modeled using both inferred and reported specifications.

Figure 24: Inferred Specification Error

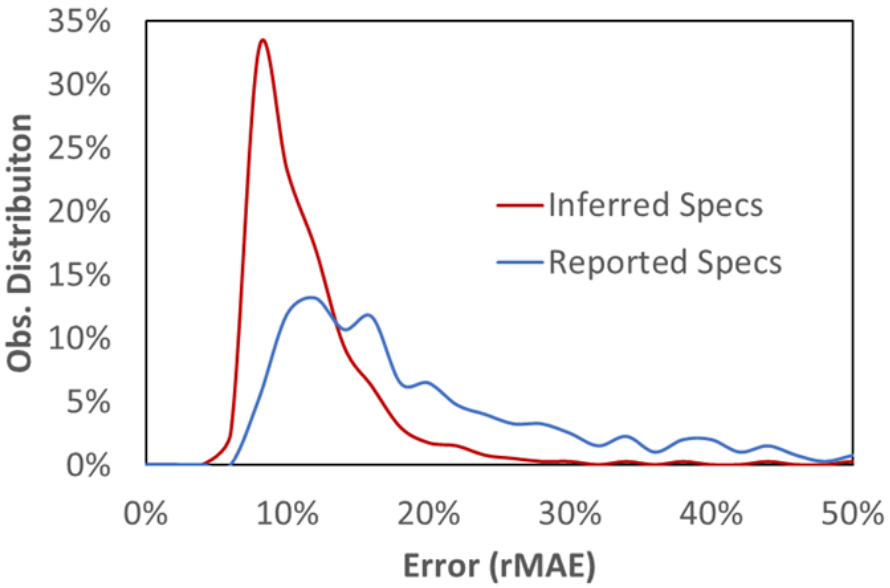

Source: Clean Power Research

Overall, inferred specifications yielded excellent results. Although production data using reported specs resulted in higher error most of the time, the results were still acceptable. There were, however, systems for which simulated data did not provide a close match with measured data. Figure 25 and Figure 27 show results for the systems with the lowest error using inferred and reported specs, respectively. Figure 26 and Figure 28 show results from systems using inferred and reported specs, respectively, which resulted in typical (median) error. In some cases, errors can reflect problems with either the measured data or the physical system itself. Degradation and soiling on some systems caused an increase in error, and although CPR worked to identify system degradation rates and soiling patterns, more research is still required before that work can lead to error reduction.

Figure 25: Inferred Specifications With Lowest Error

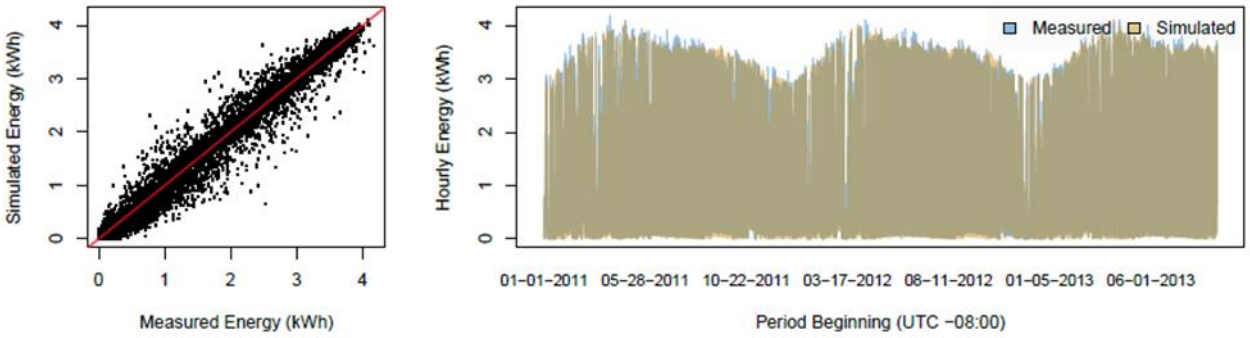

Inferred specs with lowest error (PGE-CSI-24017 at 6.3% hourly rMAE)

Source: Clean Power Research

Figure 26: Inferred Specifications With Median Error

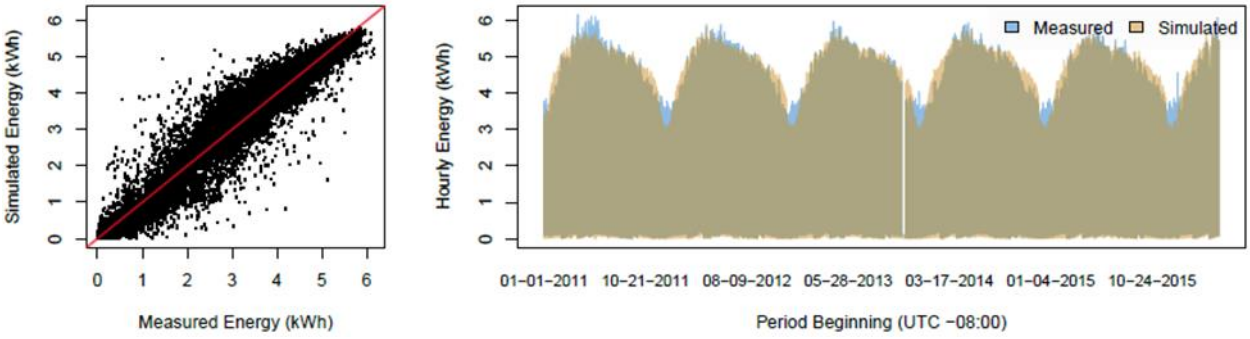

Inferred specs with median error (SCE-CSI-13299 at 10.1% hourly rMAE)

Source: Clean Power Research

Figure 27: Reported Specifications With Lowest Error

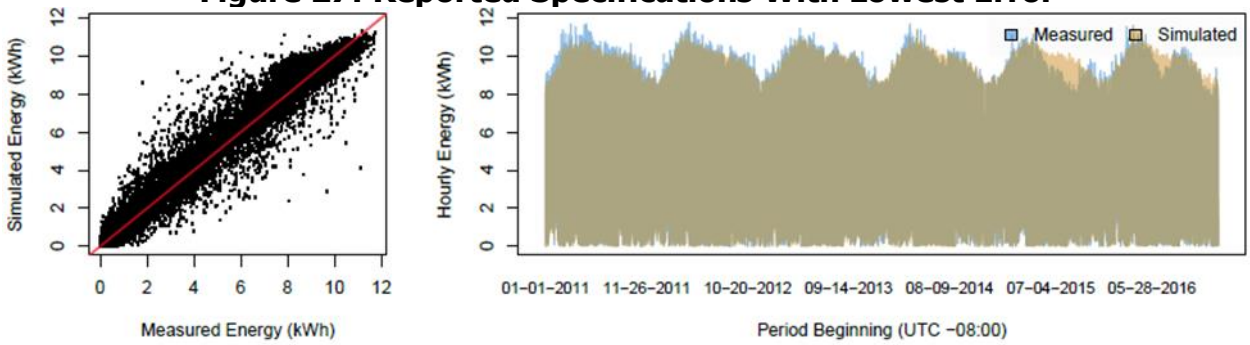

Reported specs with lowest error (SCE-CSI-09966 at 7.2% hourly rMAE)

Source: Clean Power Research

Figure 28: Reported Specifications With Median Error

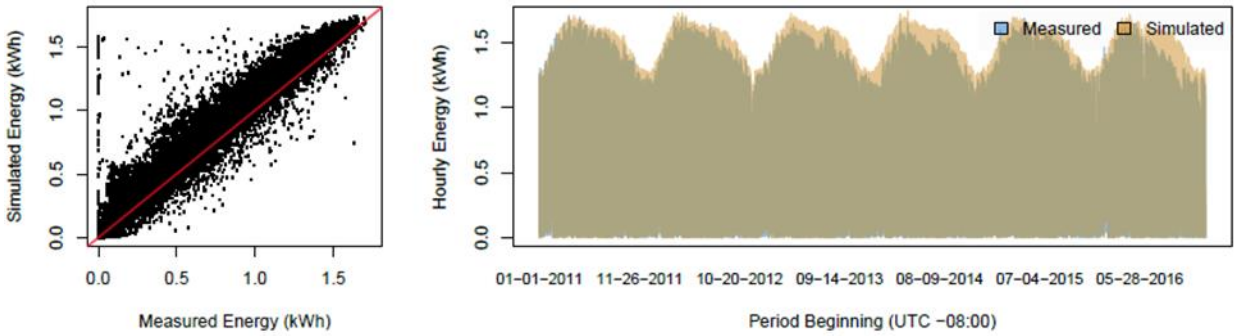

Reported specs with median error (SCE-CSI-06793 at 16.7% hourly rMAE)

Source: Clean Power Research

Conclusions and Recommendations

Although the results from the inferred system specifications can be quite good, particularly with good sets of measured data for inference, there remain several questions that deserve further exploration. For example, how much of the error in the data from the reported specifications was due to either the assumed 90 percent general derate or the lack of good data on solar obstructions? Would a blended approach, where inferred solar obstructions are used in conjunction with reported orientation, yield a much lower error? What is the effect on error of using a constant horizon with reported specifications instead of azimuth-specific solar obstructions? What techniques could be implemented to more effectively identify invalid measured data?

The DER production database, in its current form, is limited to only a single technology, solar PV, and based on only a limited sample of systems (typically one or two) within a given zip code. Because not all zip codes were included, analysis use cases were limited to those zip codes where measured data were available.

This begs the question: For the future, how can the research team design a DER production database that would broadly represent DER production profiles with other DER technologies?

To design such a database, it is useful to first consider how the database would be used. CPR proposes the following use cases for a future database:

- Developing California energy goals. Developing statewide energy goals requires an understanding of grid impacts. For example, in setting targets for electric vehicle adoption or 100-percent renewable sourcing, the database could be used to predict what the hourly load profile would look like.
- Examining differential-growth rate scenarios. Suppose load grows at a 1-percent constant rate, PV capacity grows using the logistic functions just described, and EV grows at 5 percent. The database should be able to support a study using these input assumptions.
- Defining technology requirements to meet energy objectives. Suppose that a study sought to determine how much energy storage capacity, measured in megawatts (MW) and megawatt-hours (MWh), would be needed by 2030 to maintain peak load at a

constant level? This question could be answered with appropriately scaled hourly production profiles.

- Designing incentive programs cost-effectively. Incentives could be used to target regions, step amounts, and rollout timing based on future growth scenarios; this could be supported by the database.
- Impacts of new technologies. Suppose new technologies, whether supply-side or demand-side, came into the mix. The database should be flexible enough to evaluate their impact on the electric grid.

These use cases illustrate the value of such a database, as long as it is designed to support these types of analyses. This will require separation of production by technology; for example, the PV production time series should be independently scalable from the energy storage charge/discharge time series. There should also be customer load data to determine dispatch analysis and test energy and peak demand responses. For simplicity, the database should not cover multiple years, but rather cover a single, typical year so that individual time series could be scaled to represent future-year impacts. Finally, there should be clear mapping between the database and the source data used to develop it, namely the CIDS database used to chart installations at the customer level and for metered data (similar to the Itron data used for this research).

The resulting dataset could be scaled according to analysis needs. If in year t, 500 electric vehicles (EVs) are assumed for a particular zone, then the profile for that zone's EV time series is multiplied by 500. Then, if in year t+1, 550 Evs are assumed, the same base curve is used, but the scaling factor of 550 would be used. Scaling factors would be customizable by technology, zone, and year, depending on the nature of the study.

CHAPTER 6:
Short-Term Load Forecasts

Introduction

Solar PV generation adds greater frequency of load variation than traditional, more slow-moving statistical load forecasts are able to capture. As a result, explanatory variables for estimates of solar PV generation and cloud cover movements are required to compensate for more frequent load variations. In an earlier study conducted for the California Energy Commission (CEC), *Improving Solar and Load Forecasts by Reducing Operational Uncertainty* (Publication Number CEC-500-2019-023), three alternative modeling approaches were evaluated for incorporating BTM solar PV generation into the California ISO load forecast, as described.

- Error Correction. The error correction approach relies on after-the-fact adjustments of load forecasts to account for forecasted BTM solar PV. On sunny days, the load forecast is adjusted downward, and on cloudy days the load forecast is adjusted upward. The key advantage of the error correction approach is that existing load forecast models can continue to be used unchanged. All that requires changing is a means of forecasting solar PV generation.
- Reconstituted Loads. Under the reconstituted loads approach, the historical time series of measured load is reconstituted by adding back in estimates of BTM solar PV generation. The load forecast model is then re-estimated against reconstituted loads. The subsequent reconstituted load forecasts are then adjusted after the fact by subtracting forecasts of BTM solar PV generation to accurately forecast measured loads. The advantage of this approach is that any inherent bias imposed on the estimated coefficients of a model of measured loads is controlled by estimating the model coefficients against a time series of demand for power, however it is sourced. The disadvantage is that a historical time series of BTM solar PV generation must be developed and maintained to estimate load forecast model coefficients. This approach also assumes that the historical BTM solar PV generation time series is accurate. This may not necessarily be true, in which case this approach actually places too much weight on BTM solar PV generation values.
- Direct Model. Under this approach, the weight placed on BTM solar PV generation data is estimated directly by including it as an explanatory variable in load forecast models. Weight is the estimated coefficient on the BTM solar PV generation variable. In principle, by including BTM solar PV generation as an explanatory variable, the coefficients on the remaining explanatory variables should not be biased. This approach also provides a direct forecast of measured loads that accounts for BTM solar PV generation, thus avoiding any after-the-fact processing of the load forecast. Like the reconstituted load approach, this approach requires developing and maintaining an historical time series of solar PV generation.

Based on the findings from the CEC study, the California ISO began running the reconstituted load approach alongside other production models to further improve day-ahead load

forecasting. After incorporating this approach for over a year, the California ISO concluded that it requires further refinement. Specifically, the measured load forecasts derived from the reconstituted load approach overstated the impact of solar PV generation both during the morning hours and when high temperatures created greater load demand from utility customer air conditioning.

This observation emphasizes a key assumption underlying the reconstituted load approach: that 1 kWh of BTM solar PV generation reduces measured load by 1 kWh. In effect, the estimated coefficient on the BTM solar PV generation variable equals -1.0. There are two perspectives to the question of whether this equation is true. From the BTM perspective, it is known exactly how much electricity is consumed and how much electricity was generated by solar PV panels. From the in-front-of-the-meter perspective, it is unclear whether an increase in the measured load is driven by increased consumption, by a reduction in solar PV output, or some combination of each. So measured load forecasting is essentially using a single equation to solve two unknowns.

The reconstituted load approach reduces the problem to one equation and one unknown (consumption) by assuming that the research team knew (or had a very good estimate) of the solar PV generation, and that given this estimate its impact on measured load is -1.0. By making this assumption, the problem of forecasting measured load is reduced to building a strong predictor of consumption. The California ISO observation that the reconstituted load approach appears to overstate the impact of BTM solar PV generation suggests that: some combination of: BTM solar PV generation estimates are subject to error; from the in-front-of-the-meter perspective, the impact of BTM solar PV generation is not equal to -1.0; or the consumption estimates are subject to error. This study was designed to control only the first two of those factors.

Approach

Building on the previous section that suggests that the reconstituted load approach is too restrictive, the research team developed a series of alternative model specifications that allow the impact of BTM solar PV generation to be estimated rather than imposed. The statistically adjusted BTM solar PV generation estimates are then applied to a revised reconstituted load approach to assess the forecast performance.

Four different sets of model specifications were tested. Two of the four model specifications allow the non-weather-sensitive loads to vary by day (weekday or weekend) and date range (2013 to 2015 and 2016 to 2018). The other two model specifications allow the non-weather-sensitive loads to vary by day type, date range, and season. All four specifications allow the weather-sensitive loads to vary by day type and date range. The regression model specifications use temperature curves to capture non-linear weather responses between measured loads and temperatures. The neural network model specifications use two sigmoid functions to model the non-linear weather response. The non-solar PV generation variables common to these four model specifications are described here.

Non-Seasonal Regression Model Specifications.

Non-seasonal regression model specifications have the following common set of explanatory variables:

- A linear time trend that captures long-run trends in load growth (decay).
- Day type and date range interaction terms that capture potential changes in average weekday and weekend loads in 2013-2015 and 2016-2018.
- Extremely cold and cold heating degree day spline variables with day type and date range interactions allow the estimated space-heating response to vary between weekdays and weekends and evolve before and after January 1, 2016.
- Extremely hot and hot cooling degree day spline variables with day type and date range allow the estimated space cooling response to vary between weekdays and weekends and evolve before and after January 1, 2016.

Seasonal Regression Model Specifications.

Seasonal regression model specifications have a common set of explanatory variables, including:

- A linear time trend that captures long-run trends in load growth (decay).
- Season, day type and date range interaction terms that capture potential changes in average weekday and weekend loads by season and by pre-2016 and post-2015.
- Extremely cold and cold heating degree day spline variables with day type and date range terms that allow the estimated space heating response to vary between weekdays and weekends and evolve before and after January 1, 2016.
- Extremely hot and hot cooling degree day spline variables with day type and date range interaction that allow the estimated space cooling response to vary between weekdays and weekends and evolve before and after January 1, 2016.

Non-Seasonal Neural Network Model Specifications.

The non-seasonal neural network model specifications have the following common set of explanatory variables:

- One node with a linear activation function that includes:
    - A linear time trend that captures long-run trends in load growth (decay).
    - Day type and date range interaction terms that capture potential changes in average weekday and weekend loads both pre-2016 and post-2015.
- Two nodes with a sigmoid function that includes average temperature with day type and date range interactions that allow the estimated weather response function to vary between weekdays and weekends and evolve before and after January 1, 2016.

Seasonal Neural Network Model Specifications.

The seasonal neural network model specifications have the following common set of explanatory variables:

- One node with a linear activation function that includes:
    - A linear time trend that captures long-run trends in load growth (decay).
    - Season, day type, and date range interaction terms that capture potential changes in average weekday and weekend loads both by season and pre-2016 and post-2015.

- Two nodes with a sigmoid function that include average temperatures with day type and date range that allow the estimated weather response function to vary between weekdays and weekends and evolve before and after January 1, 2016.

These base configurations for alternative model specifications build upon the goal of allowing BTM solar PV generation to vary by:

- Day type (weekdays and weekends).
- Date range (pre-2016 and post-2015).
- Temperature bin (less than 60°, 60° to 65°, 65° to 75°, and greater than 75°, all in degrees Fahrenheit) or (less than 15.6°, 15.6° to 18.3°, 18.3° to 23.9°, and greater than 23.9°, all in degrees Celsius)
- Season (winter, spring, summer, and fall)

The specific model specifications follow:

1. Base-measured load model
2. Base-reconstituted load model
3. Single variable
4. Weekday or weekend
5. Pre-2016 or post-2015
6. Weekday or weekend, pre-2016 or post-2015
7. Temperature bin
8. Temperature bin, weekday and weekend
9. Temperature bin, pre-2016 and post-2015
10. Temperature bin, weekday and weekend, pre-2016 and post-2015
11. Season
12. Season, weekday and weekend
13. Season, pre-2016 and post-2015
14. Season, weekday and weekend, pre-2016 and post-2015
15. Season, temperature bin
16. Season, temperature bin, weekday and weekend
17. Season, temperature bin, pre-2016 and post-2015
18. Season, temperature bin, weekday and weekend, pre-2016 and post-2015

Time-of-Day Estimates. Examples in the previous chapter suggest that the impact of errors in the BTM solar PV generation can vary by time of day. As a starting point for the model analysis, the research team used the 15-minute measured load data to construct loads by time of day, follows:

- Dawn = Maximum load over the time interval of 05:00 through 07:45

- Morning = Average load over the time interval of 08:00 through 10:45
- Midday = Average load over the time interval of 11:00 through 13:45
- Afternoon = Average load over the time interval of 14:00 through 16:45
- Evening = Maximum load over the time interval of 17:00 through 19:45

Maximum load is used for the dawn and evening time periods to allow for the greatest possible impact of BTM solar PV generation variables. This is particularly useful during the fall and winter seasons when BTM solar PV generation data goes to zero for a number of time intervals. The full set of model specifications are estimated separately for each of the five time-of-day periods.

Results

The perceived impact of BTM solar PV generation on measured load can vary by time of day, season, and temperature. The following analysis is for SDG&E. Comparable analyses were also done for PG&E and SCE.

As a starting point for the SDG&E analysis, the 15-minute level measured load data were used to construct measured loads by time of day, as follows:

- Dawn = Maximum load over the time interval of 05:00 through 07:45
- Morning = Average load over the time interval of 08:00 through 10:45
- Midday = Average load over the time interval of 11:00 through 13:45
- Afternoon = Average load over the time interval of 14:00 through 16:45
- Evening = Maximum load over the time interval of 17:00 through 19:45

Maximum load is used for the dawn and evening time periods to allow the greatest possible impact of the BTM solar PV generation variables. This is particularly useful during the fall and winter seasons when the BTM solar PV generation data goes to zero for a number of these time intervals.

Measured Load and Behind-the-Meter Solar Photovoltaic Generation Data

Measured load, BTM solar PV generation, temperature, dew point, wind speed, and cloud cover data are available for January 1, 2013 through May 7, 2018.

Only midday data and results are shown. The measured load and corresponding BTM solar PV generation data for midday intervals appear in Figure 29 and Figure 30. Observations about these data include that:

- BTM solar PV generation has grown steadily over this period.
- There is a noticeable step-up in the BTM solar PV generation starting near the beginning of 2016. This is most noticeable when viewing the dawn BTM solar PV generation data.
- There is a noticeable downward trend in measured loads starting around 2016. This is consistent with the increase in BTM solar PV generation.

Figure 29: San Diego Gas & Electric Measured Load – Midday

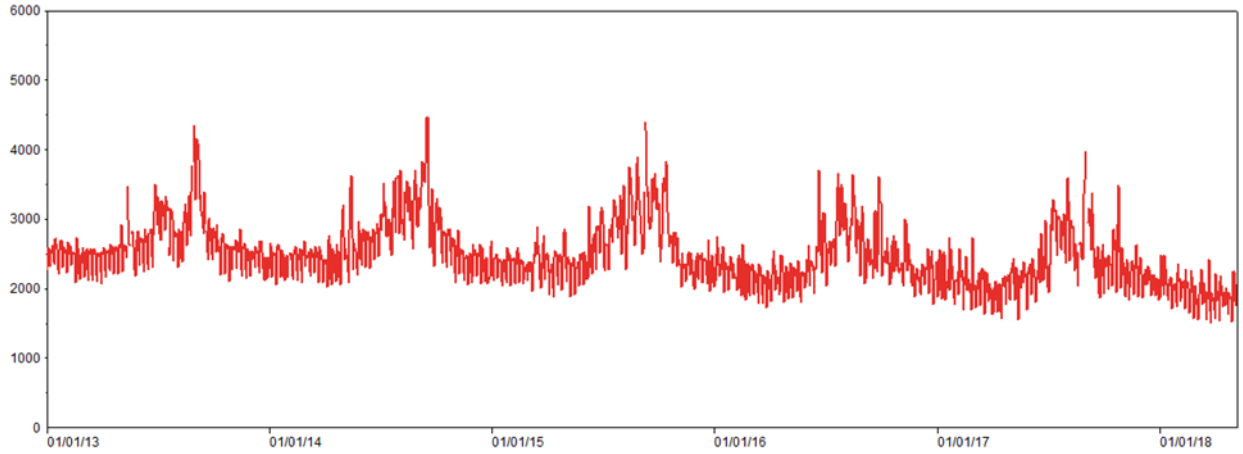

Source: Clean Power Research

Figure 30: San Diego Gas & Electric Solar Photovoltaic Generation – Midday

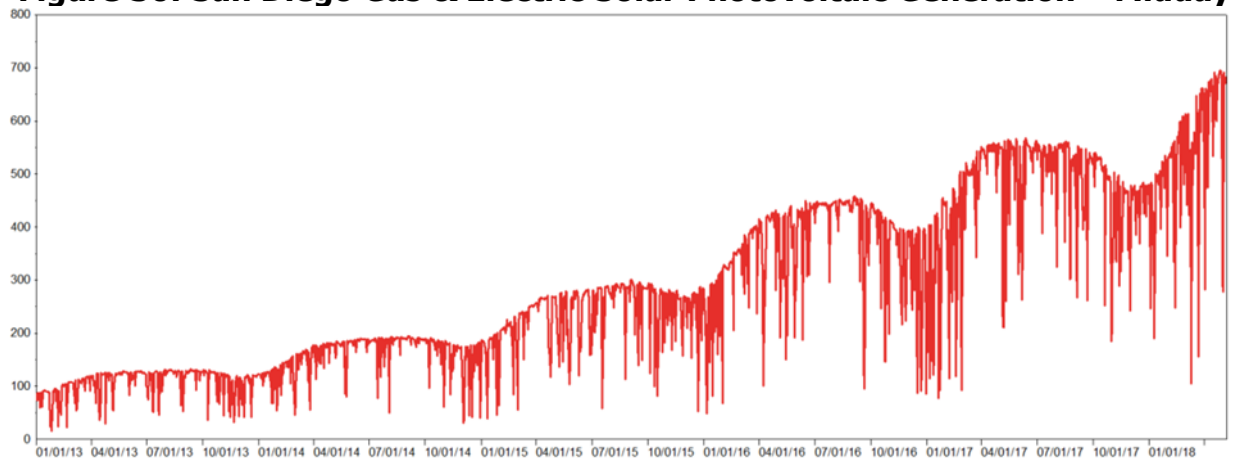

Source: Clean Power Research

Weather Response

Figure 31 and Figure 32 show the measured load response to changes in temperatures. Again, only the midday interval is shown.

In these charts, pre-2016 weekday loads are blue, pre-2016 weekend day loads are purple, post-2015 weekday loads are green, and post-2015 weekend day loads are gold. For each time-of-day period, two charts are presented. The first is the observed measured load weather response. The second is the estimated measured load weather response. Estimated weather response functions are based on a three-node neural network with the following configuration.

- Node 1 has a linear activation function and includes pre-2016 weekday and weekend binary variables and post-2015 weekday and weekend binary variables.
- Node 2 has a sigmoid function and includes the time-of-day average temperature, the pre-2016 weekday and weekend binary variables, and post-2015 weekday and weekend binary variables.

- Node 3 has a sigmoid function and includes the time-of-day average temperature, the pre-2016 weekday and weekend binary variables, and post-2015 weekday and weekend binary variables.

The binary variables on the linear node (Node 1) allow the non-weather sensitive portion of measured load to vary across weekdays and weekend days, as well as across the pre-2015 and post-2016 time periods.

The two nonlinear nodes (Node 2 and Node 3) work together to approximate the nonlinear weather response. The binary variables on the nonlinear nodes (Node 2 and Node 3) allow the weather response to differ across weekdays and weekends, as well as across the pre-2015 and post-2016 periods.

Observations about these charts follow.

- The non-weather sensitive portion of the time-of-period load shows a reduction between the pre-2015 and post-2016 periods. The reduction in non-weather sensitive loads is most pronounced in the morning, midday, and afternoon periods, which is when BTM solar PV generation is highest. This suggests that the penetration of BTM solar PV over this time is a factor in the downward trend in loads in the daily graphs of time-of-period loads.
- The reduction in the non-weather sensitive portion of the time-of-period load appears to be consistent across weekdays and weekend days. This further supports the idea that the penetration of BTM solar PV generation is a factor in the downward trend in measured loads.
- Focusing on the dawn estimated weather response function, there is no noticeable change in the weather response over time. The reduction in the response function post-2015 appears to be at all temperatures.
- Focusing on the morning estimated weather response function, there is a noticeable reduction in the weather response for both weekdays and weekend days post-2015. The post-2015 weekday weather response appears to be steeper on the space-heating side and flatter on the air-conditioning side. The weekend day air-conditioning response appears to be flatter post-2015.
- Focusing on the midday estimated weather response function, the post-2015 weekday weather response is close to the pre-2016 weekend day response. The gap between the pre-2016 weekday response and the post-2015 weekday response is greatest when temperatures are in the 70s, but reduces quickly as temperatures cool. The gap reduces, but not as quickly when temperatures go above 80. The gap between the pre-2016 weekend day response and the post-2015 weekend day response appears constant at all temperatures.
- Focusing on the afternoon estimated weather response function, the post-2015 weekday response is close to the pre-2016 weekend day response. The gap between the pre-2016 weekday response and the post-2015 weekday response is greatest when temperatures are in the 70s but reduces quickly as temperatures either cool or increase. In a similar fashion, the gap between the pre-2016 weekend day response and the post-2015 weekend day response is greatest when temperatures are in the 70s but reduces quickly as temperatures either cool or increase.

- Focusing on the evening estimated weather response function, there is a slight change in the level of weather response over time. The reduction in the weekday response function appears to be at all temperatures. There is a greater reduction in the post-2015 weekend weather response when temperatures fall below 70 degrees than when temperatures rise above 70 degrees.

Figure 31: San Diego Gas & Electric Observed Measured Load Weather Response – Midday

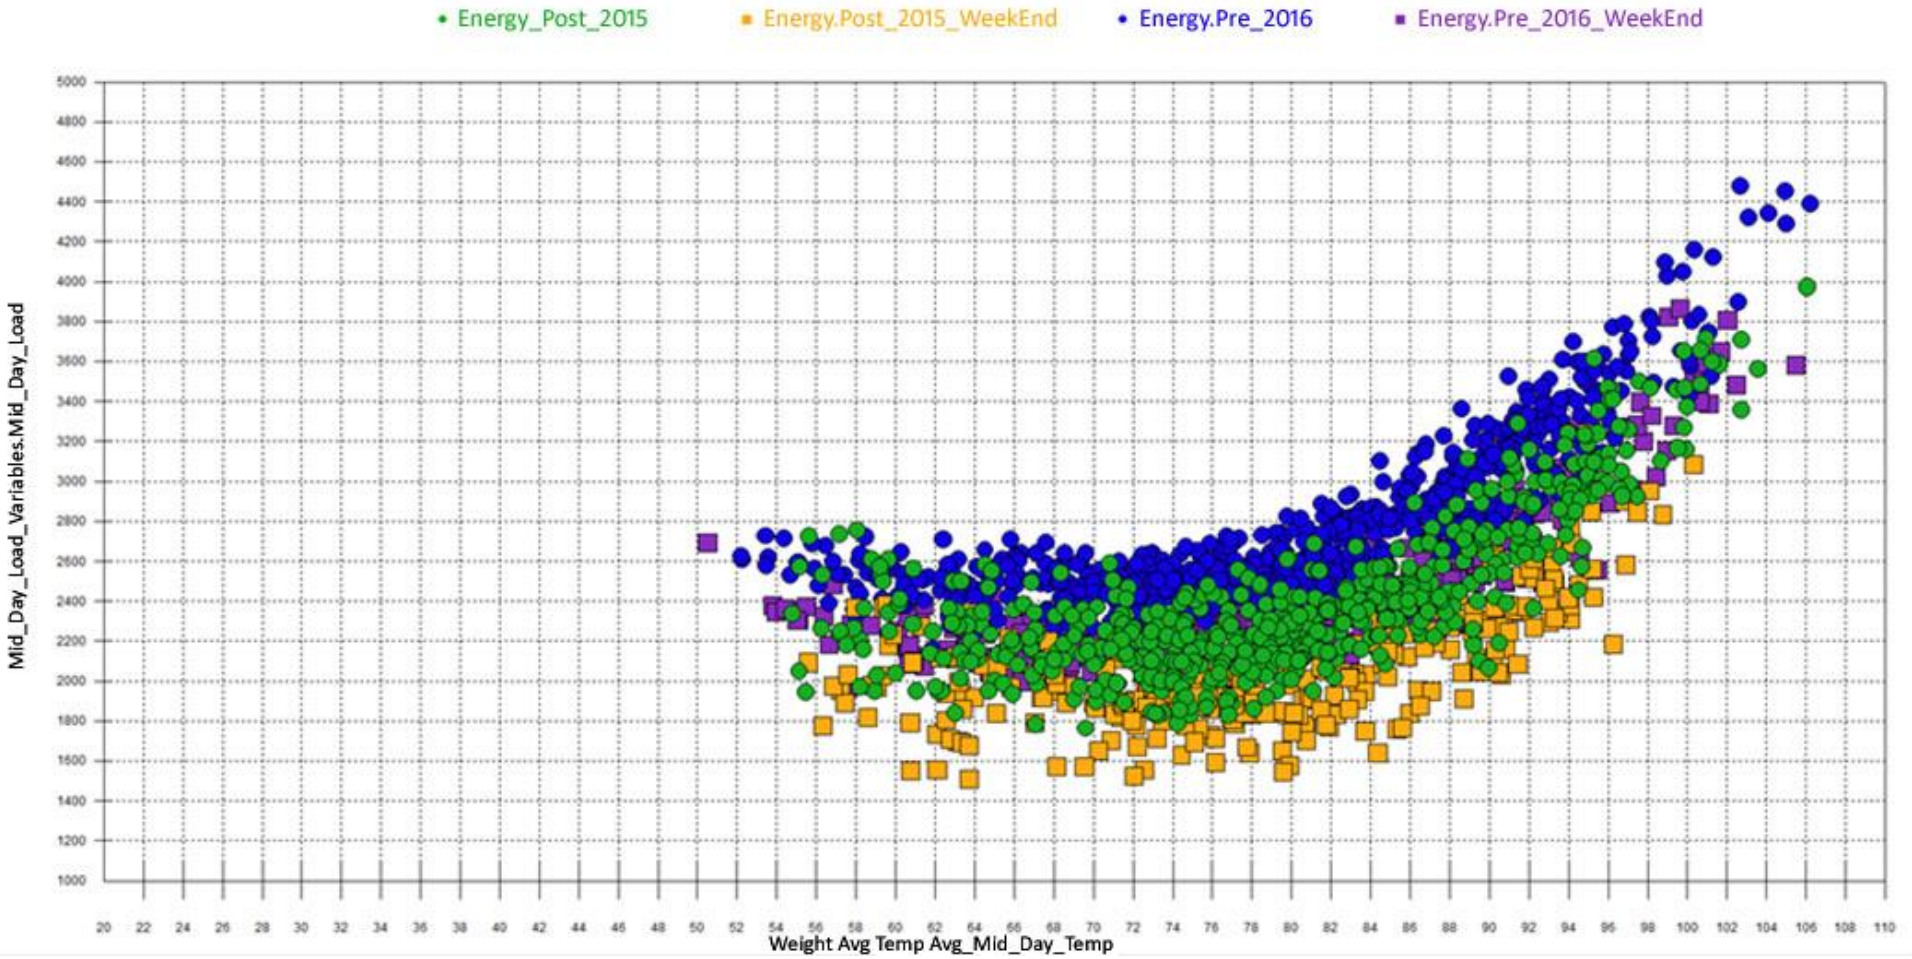

Source: Clean Power Research

Figure 32: San Diego Gas & Electric Neural Network Estimated Measured Load Weather Response – Midday

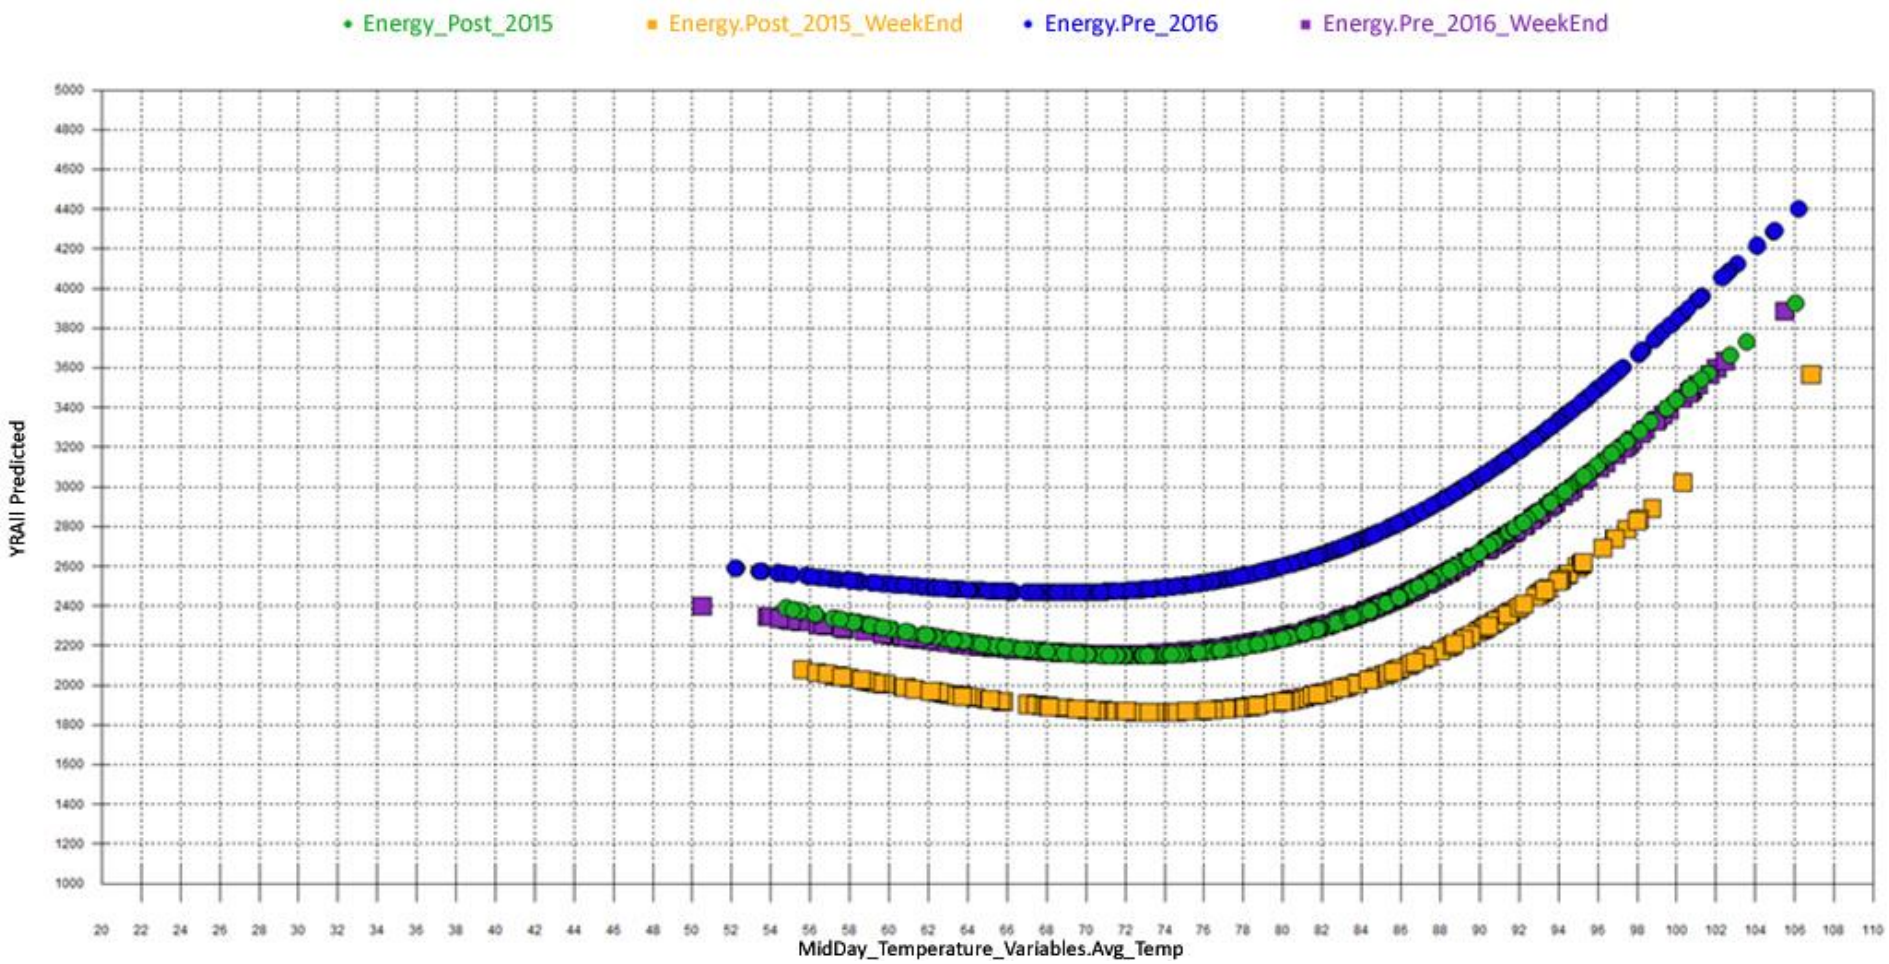

Source: Clean Power Research

San Diego Gas & Electric Model Result Comparison

This section presents the results of the alternative BTM solar PV generation model specifications just described. The results are presented as a cascade of models; each step in the cascade allows the impact of the BTM solar PV generation to vary by day type, date range, temperature bin, and season.

Table 1 presents the estimated aggregate BTM solar PV generation coefficients by model specification. For reference, the assumed coefficient of -1.0 is included in this table. The columns represent results from the neural network (NNET) and regression (REG) model specifications. To ease comparison across the alternative model specifications, only the estimated aggregate coefficients are presented. The estimated aggregate coefficients are computed as the weighted average of the estimated coefficients across all BTM solar PV generation variables included in the model specification. For example, the estimated coefficients for the dawn regression model with pre-2016 and post-2015 interactions are:

- Solar PV generation * pre-2016 = -0.496
- Solar PV generation * post-2015 = -0.488
- The in-sample average solar PV generation pre-2016 data = 38.980
- The in-sample average solar PV generation post-2015 data = 63.842

The estimated aggregate solar PV generation coefficient is computed as:

- Aggregate BTM solar PV coefficient = [(-0.496 * 38.980) + (-0.488 * 63.842)]/[38.980 + 63.842]
- Aggregate BTM solar PV coefficient = -0.491

Observations about the estimated aggregate coefficients appear here.

- Dawn Hours. The estimated aggregate BTM solar PV adjustment range is between -0.664 to 0.366. The positive values indicate the difficulty of isolating the impact of BTM solar PV generation from other factors that swing measured load. They also reflect the model rather than their specifications, especially when all interaction terms are included. Focusing on the negative estimated coefficients suggests that the assumption of -1.0 is too restrictive in the dawn hours. The average estimated impact without seasonal terms is -0.55 and -0.24 with the seasonal terms. Non-seasonal estimates suggest almost halving the BTM solar PV generation estimates. The less than -1.0 coefficient estimates might reflect errors with the BTM solar PV generation estimates, which could be driven by the azimuth, tilt, soiling, and shading assumptions underlying the estimates. Alternatively, because BTM solar PV generation is small relative to the average dawn load it might be difficult for the statistical models to isolate the BTM solar PV impact from other factors driving dawn loads. The fact that most of the models lead to an estimated negative impact less than -1.0 suggests that adjusting the BTM solar PV estimates is warranted. When the season binary variables are introduced to the model specification the estimated BTM solar PV impact is reduced. This suggests that some of the BTM solar PV generation impact correlates with seasonal load variations. In summary, the results suggest using a statistical adjustment factor of around -0.24 for dawn hours.

- Morning Hours. The estimated aggregate BTM solar PV adjustment range is between -1.01 to -1.25. The average estimated impact without the seasonal terms is -1.19 and -1.10 with the seasonal terms. The estimated coefficients suggest that the assumption of -1.0 is too restrictive. Using the results from the seasonal model specifications suggests a statistical adjustment factor of around -1.10 for the morning hours.
- Midday Hours. The estimated aggregate BTM solar PV adjustment range is between -1.28 to -1.55. The average estimated impact without the seasonal terms is -1.38 and -1.37 with the seasonal terms. The estimated coefficients suggest that the assumption of -1.0 is too restrictive. Using the results from the seasonal model specifications suggests a statistical adjustment factor of around -1.37 for the midday hours.
- Afternoon Hours. The estimated aggregate BTM solar PV adjustment range is between 0.00 to -1.63. The 0.00 estimated coefficient reflects model over specification when all interaction terms are included. The average estimated impact without the seasonal terms is -1.57, and -1.35 with the seasonal terms. The estimated coefficients suggest that the assumption of -1.0 is too restrictive. Using the results from the seasonal model specifications suggests a statistical adjustment factor of around -1.35 for the afternoon hours.
- Evening Hours. The estimated aggregate BTM solar PV adjustment range is between -2.18 to 1.22. These positive values reflect model over specification especially when all interaction terms are included. The average estimated impact without the seasonal terms is -1.95, and -1.23 with the seasonal terms. The estimated coefficients suggest that the assumption of -1.0 is too restrictive. Using the results from the seasonal model specifications suggests a statistical adjustment factor of around -1.23 for the evening hours.

Table 1: Estimated Aggregate Solar Photovoltaic Generation Coefficient

	Dawn		Morning		Midday		Afternoon		Evening	
	NNET	REG	NNET	REG	NNET	REG	NNET	REG	NNET	REG
Reconstituted Load	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.000
Single Variable	-0.634	-0.642	-1.133	-1.144	-1.326	-1.368	-1.538	-1.598	-1.782	-1.846
Week Day/Weekend	-0.634	0.342	-1.145	-1.140	-1.323	-1.369	-1.541	-1.603	-1.783	-1.846
Pre-2016/Post 2015	-0.659	-0.664	-1.180	-1.198	-1.338	-1.395	-1.510	-1.604	-2.042	-2.175
Week Day/Weekend, Pre-2016/Post-2015	-0.657	-0.663	-1.174	-1.195	-1.343	-1.395	-1.513	-1.609	-2.070	-2.174
Temperature Bin	-0.553	-0.551	-1.196	-1.219	-1.345	-1.450	-1.523	-1.621	-1.725	-1.866
Temperature Bin, Week Day/Weekend	-0.555	-0.552	-1.167	-1.212	-1.357	-1.450	-1.529	-1.626	-1.735	-1.866
Temperature Bin, Pre-2016/Post-2015	-0.599	-0.604	-1.210	-1.253	-1.356	-1.454	-1.498	-1.622	-1.979	-2.186
Temperature Bin, Pre-2016/Post-2015, Week Day/Weekend	-0.598	-0.603	-1.205	-1.248	-1.381	-1.452	-1.503	-1.626	-1.980	-2.184
Season	-0.276	-0.265	-1.027	-1.041	-1.282	-1.341	-1.340	-1.497	-1.372	-1.695
Season, Week Day/Weekend	-0.272	-0.264	-1.010	-1.031	-1.288	-1.340	-1.346	-1.502	-1.415	-1.704
Season, Pre-2016/Post-2015	-0.263	-0.247	-1.053	-1.078	-1.302	-1.405	-1.339	-1.556	-1.509	-1.848

	Dawn		Morning		Midday		Afternoon		Evening	
	NNET	**REG**	**NNET**	**REG**	**NNET**	**REG**	**NNET**	**REG**	**NNET**	**REG**
Season, Week Day/Weekend, Pre-2016/Post 2015	-0.258	-0.246	-1.036	-1.067	-1.301	-1.403	-1.345	-1.561	-1.519	-1.855
Season, Temperature Bin	-0.283	-0.301	-1.130	-1.138	-1.287	-1.390	-1.315	-1.450	-1.411	-1.678
Season, Temperature Bin, Week Day/Weekend	-0.286	-0.307	-1.129	-1.129	-1.286	-1.389	-1.326	-1.458	0.499	-1.690
Season, Temperature Bin, Pre-2016/Post 2015	-0.291	-0.300	-1.177	-1.224	-1.391	-1.510	-1.410	-1.579	1.216	-1.871
Season, Temperature Bin, Week Day/Weekend, Pre- 2015/Post 2015	0.366	-0.304	-1.120	-1.224	-1.550	-1.460	0.000	-1.540	0.000	-1.822

Source: Clean Power Research

To determine the potential for forecast improvements, the statistically adjusted BTM solar PV estimates derived from each model specification were used in the reconstituted load approach. The resulting in-sample mean absolute percentage errors (MAPEs) in these statistically adjusted BTM solar PV generation estimates are then compared with the original reconstituted load approach using the unadjusted BTM solar PV generation estimates. For comparison purposes, the corresponding direct model results are presented in Table 2 through Table 5. The cells highlighted in green represent the best in-sample MAPE for that time period and group of model specifications. The cells highlighted in red represent the best in-sample MAPE for that time period across all model specifications. Observations about these results follow.

- In general, the reconstituted load approach is enhanced by using statistically adjusted BTM solar PV generation estimates.
- The direct model approach yields, on average, better in-sample fits than the reconstituted load approach.
- Except for the evening time-of-use period, model specifications allowing the BTM solar PV generation impact to vary by temperature bin yield the best in-sample fits.
- The evening time-of-use period favors a model specification that allows the BTM solar PV generation impact to vary between weekdays and weekend days.
- There is slight evidence that the BTM solar PV generation impact has evolved over the estimation range.
- Focusing on results for the seasonal NNET, the morning hours have the greatest improvement over the baseline reconstituted load approach, with a reduction from an in-sample MAPE of 5.10 percent to 3.04 percent MAPE.
- The dawn hours benefit from the statistically adjusted BTM solar PV estimate with an improvement in the in-sample MAPE from 2.99 percent to 2.72 percent.
- The midday hours show marginal improvements, with the in-sample MAPE improving from 3.05 percent to 2.93 percent.
- The afternoon hours show a 14th of 1 percent improvement, with the in-sample MAPE reducing from 2.96 percent to 2.82 percent.
- The evening hours also show improvement, with the in-sample MAPE reducing from 3.36 percent to 2.82 percent.

Table 2: Revised Reconstituted Load Versus Direct Model: Non-Seasonal Regression

	Dawn		Morning		Mid-Day		Afternoon		Evening	
Regression Model	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
No SolarPVGen Treatment		3.46%		4.18%		5.02%		5.22%		4.82%
Reconstituted Load with 100% SolarPVGen	3.28%		3.57%		4.11%		4.40%		4.55%	
Single Variable	3.22%	3.21%	3.56%	3.55%	4.02%	4.00%	4.21%	4.19%	4.13%	4.11%
Week Day/Weekend	3.73%	3.06%	3.56%	3.55%	4.02%	4.00%	4.20%	4.18%	4.13%	4.11%
Pre-2016/Post 2015	3.21%	3.21%	3.55%	3.54%	4.02%	4.00%	4.21%	4.19%	3.92%	3.90%
Week Day/Weekend, Pre-2016/Post-2015	3.20%	3.20%	3.55%	3.54%	4.02%	4.00%	4.20%	4.18%	3.91%	3.89%
Temperature Bin	3.16%	3.15%	3.53%	3.52%	3.87%	3.85%	4.10%	4.08%	4.10%	4.08%

	Dawn		Morning		Mid-Day		Afternoon		Evening	
Regression Model	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
Temperature Bin, Week Day/Weekend	3.15%	3.15%	3.53%	3.52%	3.87%	3.85%	4.09%	4.07%	4.09%	4.08%
Temperature Bin, Pre-2016/Post-2015	3.14%	3.13%	3.52%	3.51%	3.84%	3.82%	4.06%	4.04%	3.88%	3.86%
Temperature Bin, Pre-2016/Post-2015, Week Day/Weekend	3.12%	3.12%	3.57%	3.50%	3.90%	3.81%	4.04%	4.03%	4.02%	3.85%
Best MAPE	**3.06%**		**3.50%**		**3.81%**		**4.03%**		**3.85%**	

The cells highlighted in green represent the best in-sample MAPE for that time-period and group of model specifications. The cells highlighted in red represent the best in-sample MAPE for that time-period across all model specifications.

Source: Clean Power Research

Table 3: Revised Reconstituted Load Versus Direct Model: Seasonal Regression

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct
No SolarPVGen Treatment		3.06%		3.42%		3.94%		4.10%		4.10%
Reconstituted Load with 100% SolarPVGen	3.16%		3.37%		3.98%		4.14%		4.13%	
Single Variable	2.98%	2.97%	3.37%	3.35%	3.92%	3.90%	4.04%	4.03%	3.73%	3.71%
Season	2.95%	2.95%	3.36%	3.34%	3.92%	3.90%	4.03%	4.02%	3.73%	3.71%
Week Day/Weekend	2.94%	2.94%	3.35%	3.33%	3.92%	3.90%	4.03%	4.01%	3.73%	3.71%
Pre-2016/Post 2015	2.94%	2.86%	3.34%	3.23%	3.92%	3.66%	4.03%	3.74%	3.69%	3.59%
Week Day/Weekend, Pre-2016/Post-2015	2.93%	2.93%	3.33%	3.31%	3.91%	3.89%	4.03%	4.01%	3.68%	3.66%
Temperature Bin	2.90%	2.90%	3.29%	3.27%	3.75%	3.73%	3.82%	3.80%	3.68%	3.65%

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
Temperature Bin, Week Day/Weekend	2.88%	2.88%	3.27%	3.26%	3.75%	3.73%	3.81%	3.79%	3.66%	3.64%
Temperature Bin, Pre-2016/Post-2015	3.00%	2.86%	3.67%	3.23%	4.30%	3.66%	4.17%	3.74%	3.67%	3.59%
Temperature Bin, Pre-2016/Post-2015, Week Day/Weekend	2.83%	2.83%	3.22%	3.20%	3.66%	3.64%	3.73%	3.71%	3.59%	3.57%
Best MAPE	**2.83%**		**3.20%**		**3.64%**		**3.71%**		**3.57%**	

The cells highlighted in Green represent the best in-sample MAPE for that time-period and group of model specifications. The cells highlighted in Red represent the best in-sample MAPE for that time-period across all model specifications.

Source: Clean Power Research

Table 4: Revised Reconstituted Load Versus Direct Model: Non-Seasonal Neural Network

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
No SolarPVGen Treatment		3.28%		4.07%		4.34%		4.39%		4.19%
Reconstituted Load with 100% SolarPVGen	3.12%		3.45%		3.29%		3.39%		3.76%	
Single Variable	3.05%	3.04%	3.44%	3.44%	3.18%	3.17%	3.15%	3.15%	3.20%	3.19%
Week Day/Weekend	3.05%	3.03%	3.44%	3.46%	3.18%	3.17%	3.14%	3.14%	3.20%	3.19%
Pre-2016/Post 2015	3.05%	3.03%	3.43%	3.43%	3.17%	3.17%	3.16%	3.15%	3.05%	3.05%
Week Day/Weekend, Pre-2016/Post-2015	3.04%	3.02%	3.43%	3.43%	3.18%	3.17%	3.15%	3.14%	3.04%	3.04%
Temperature Bin	3.03%	3.01%	3.43%	3.44%	3.18%	3.17%	3.14%	3.14%	3.20%	3.19%

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
Temperature Bin, Week Day/Weekend	3.01%	3.00%	3.43%	3.43%	3.18%	3.17%	3.14%	3.13%	3.19%	3.17%
Temperature Bin, Pre-2016/Post-2015	3.01%	3.00%	3.42%	3.41%	3.17%	3.16%	3.15%	3.15%	2.97%	2.95%
Temperature Bin, Pre-2016/Post-2015, Week Day/Weekend	3.13%	2.98%	3.54%	3.40%	3.36%	3.16%	3.19%	3.12%	3.17%	2.94%
Best MAPE	**2.98%**		**3.40%**		**3.16%**		**3.12%**		**2.94%**	

The cells highlighted in green represent the best in-sample MAPE for that time-period and group of model specifications. The cells highlighted in red represent the best in-sample MAPE for that time-period across all model specifications

Source: Clean Power Research

Table 5: Revised Reconstituted Load Versus Direct Model: Seasonal Neural Network

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct	Recon-stituted	Direct
No SolarPVGen Treatment		2.90%		3.24%		3.05%		2.95%		3.17%
Reconstituted Load with 100% SolarPVGen	2.99%		5.10%		3.05%		2.96%		3.36%	
Single Variable	3.00%	2.80%	3.42%	3.16%	3.41%	2.99%	3.14%	2.89%	3.00%	2.90%
Season	2.78%	2.78%	3.16%	3.15%	2.97%	2.98%	2.86%	2.87%	2.89%	2.99%
Week Day/Weekend	2.77%	2.76%	3.15%	3.15%	2.98%	2.97%	2.87%	2.86%	2.89%	2.88%
Pre-2016/Post 2015	2.77%	2.76%	3.14%	3.14%	2.97%	2.97%	2.88%	2.87%	2.85%	2.83%
Week Day/Weekend, Pre-2016/Post-2015	2.76%	2.75%	3.13%	3.12%	2.96%	2.95%	2.87%	2.86%	2.84%	2.82%
Temperature Bin	2.75%	2.74%	3.08%	3.07%	2.96%	2.96%	2.84%	2.84%	2.87%	2.86%

	Dawn		Morning		Midday		Afternoon		Evening	
	Recon-stituted	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**	**Recon-stituted**	**Direct**
Temperature Bin, Week Day/Weekend	2.73%	2.72%	3.05%	3.05%	2.95%	2.95%	2.84%	2.83%	13.10%	31.18%
Temperature Bin, Pre-2016/Post-2015	2.82%	2.72%	3.54%	3.04%	3.72%	2.93%	3.30%	2.82%	14.93%	34.93%
Temperature Bin, Pre-2016/Post-2015, Week Day/Weekend	12.12%	10.53%	3.14%	3.57%	28.39%	62.20%	43.80%	4.28%	29.72%	6.93%
Best MAPE	**2.72%**		**3.04%**		**2.93%**		**2.82%**		**2.82%**	

The cells highlighted in green represent the best in-sample MAPE for that time-period and group of model specifications. The cells highlighted in red represent the best in-sample MAPE for that time-period across all model specifications

Source: Clean Power Research

Conclusions and Recommendations

Study results suggest that the current reconstituted load approach should be augmented to allow the assumed impact of BTM solar PV generation on loads to vary by temperature bin, season, weekday versus weekend day, and pre-2016/post-2015. As a practical matter, a complete set of season, temperature bin, weekday/weekend and pre-2016/post-2015 interaction terms leads to model over specification. To control for over specification, the research team recommends pulling back from model specifications with season interaction terms. This reduces the list of candidate-model specifications to those that allow solar PV generation impact to vary by: temperature bin; temperature bin and weekday/weekend; and temperature bin, weekday/weekend, and pre-2016/post-2015.

Conclusions and Recommendations

This study's goal was to improve load forecasting accuracy of the California ISO's existing reconstituted load forecast models by undertaking several initiatives.

Initiative 1. Improve the Existing Reconstituted Load Forecast Framework

Findings suggest that the assumption that 1 kWh of BTM solar PV generation leads to a 1 kWh reduction in measured load is too restrictive. To allow greater accuracy of the reconstituted load forecast, results suggest that BTM solar PV impacts should be allowed to vary by utility, time-of-day, and temperature range.

General conclusions about these data include:

- Dawn. Over the model estimation period the ratio of Dawn solar PV generation to Dawn load averaged 5.06% for SDG&E, 3.40% for PG&E, and 3.09% for SCE. The estimated coefficients for SDG&E and SCE suggest reducing the impact of BTM solar PV generation by about 40%, from -1.0 to -0.55 for SDG&E and -0.57 for SCE. In contrast, the positive estimated coefficients for PG&E suggest the penetration of solar PV was not enough over the model estimation period to separate the impact from a general swing up in loads during the Spring and Summer season. There is some evidence that this will change for PG&E as solar PV systems continues to be installed.
- Morning. Over the model estimation period the ratio of Morning solar PV generation to Morning load averaged 9.41% for SDG&E, 6.71% for PG&E, and 5.38% for SCE. For SDG&E, the net effect for the Morning hours is to lift apparent impact of the BTM solar PV generation estimates by roughly 20%. For SCE, the net effect is to lift apparent impact the BTM solar PV generation estimates by about 40%. For PG&E, the net effect is to lower the apparent impact of the BTM solar PV generation estimates by about 30%.
- Midday. Over the model estimation period the ratio of midday solar PV generation to Mid-Day load averaged 12.30% for SDG&E, 9.12% for PG&E, and 6.55% for SCE. For SDG&E, the net effect for the midday hours is to lift apparent impact of the BTM solar PV generation estimates by roughly 45%. For SCE, the net effect is to lift apparent impact the BTM solar PV generation estimates by about 250%. For PG&E, the net effect is to life the apparent impact of the BTM solar PV generation estimates by about 50%.
- Afternoon. Over the model estimation period the ratio of Afternoon solar PV generation to Afternoon load averaged 7.31% for SDG&E, 6.02% for PG&E, and 3.96% for SCE.

For SDG&E, the net effect for the Afternoon hours is to lift apparent impact of the BTM solar PV generation estimates by roughly 60%. For SCE, the net effect is to lift apparent impact the BTM solar PV generation estimates by about 300%. For PG&E, the net effect is to life the apparent impact of the BTM solar PV generation estimates by about 50%.

- Evening. Over the model estimation period the ratio of Evening solar PV generation to Evening load averaged 2.16% for SDG&E, 2.36% for PG&E, and 1.35% for SCE. For SDG&E, the net effect for the Afternoon hours is to lift apparent impact of the BTM solar PV generation estimates by roughly 90%. For SCE, the net effect is to lift apparent impact the BTM solar PV generation estimates by about 320%. For PG&E, the net effect is to lower the apparent impact of the BTM solar PV generation estimates by about 40%.

Initiative 2. Improved Estimated Weather Response

This initiative determined whether the response of measured loads to weather variations has evolved with the increasing penetration of BTM solar PV installations. A series of statistical models were therefore used to estimate the weather response function for the years before 2016 and after 2015. A comparison of the estimated weather response functions was used to determine whether the overall weather response has evolved over the last five years.

The findings from this initiative are:

San Diego Gas & Electric Findings

- Focusing on the dawn estimated weather response function, there is no noticeable change in the weather response over time. The reduction in the response function post-2015 appears to be at all temperatures.
- Focusing on the morning estimated weather response function, there is a noticeable lowering in the weather response for both weekday and weekend days post-2015. The post-2015 weekday weather response appears to be steeper on the space-heating side and flatter on the air-conditioning side. The weekend day air-conditioning response appears to be flatter post-2015.
- Focusing on the midday estimated weather response function, the post-2015 weekday weather response is close to the pre-2016 weekend day response. The gap between the pre-2016 weekday response and the post-2015 weekday response is greatest when temperatures are in the 70s, but reduces quickly as temperatures cool. The gap reduces, but not as quickly, when temperatures rise above 80 degrees. The gap between the pre-2016 weekend day response and the post-2015 weekend day response appears to be constant at all temperatures.
- Focusing on the afternoon estimated weather response function, the post-2015 weekday response is close to the pre-2016 weekend day response. The gap between the pre-2016 weekday response and the post-2015 weekday response is greatest when temperatures are in the 70s but reduces quickly as temperatures grow either colder or hotter. In a similar fashion, the gap between the pre-2016 weekend day response and the post-2015 weekend day response is greatest when temperatures are in the 70s, but reduces quickly as temperatures grow either colder or hotter.

- Focusing on the evening estimated weather response function, there is slight change in the level of the weather response over time. The reduction in the weekday response function appears to be at all temperatures. There is a greater reduction in the post-2015 weekend weather response when temperatures fall below 70 degrees than when temperatures rise above 70 degrees.

Pacific Gas and Electric Findings

- The overall dawn weather response shows a distinct weekday versus weekend response. This is expected given the strong mix of residential and small commercial loads found in PG&E's territory.
- Unlike other time-of-day periods, there is only a slight change in the observed and estimated dawn weather response pre-2016 and post-2015.
- There is an apparent downward shift in the non-weather-sensitive portion of morning loads between the pre-2015 and post-2016 periods. This is apparent in both the actual and estimated morning weather response functions.
- Based on the estimated morning weather response function, there is a steepening in the space-heating response for both weekdays and weekends post-2015. There is a strong flattening in the weekend space-cooling response post-2015. There is no obvious change in the weekday space cooling response post-2015.
- On average, estimated midday weekday and weekend weather response functions are lower post-2015. The estimated weekday response is roughly higher than the estimated pre-2016 weekend response.
- The midday space cooling response appears roughly unchanged post-2015. In contrast, the midday space heating response appears steeper on weekdays and flatter on weekend days.
- On average, the estimated afternoon weekday and weekend weather response functions are lower post-2015. The estimated weekday response is roughly higher than the estimated pre-2016 weekend response.
- The estimated afternoon weekday space-cooling response appears steeper post-2015. The estimated weekend space-cooling response appears flatter post-2015. Both the estimated afternoon weekday and weekend day space-heating responses appear steeper post-2015.
- On average, the estimated evening weekday and weekend weather response functions are lower post-2015.
- The estimated evening weekday space-cooling response appears steeper post-2015. The estimated weekend space-cooling response appears roughly unchanged post-2015. Both the estimated evening weekday and weekend day space-heating responses appear flatter post-2015.

Southern California Edison Findings

- Focusing on the dawn estimated weather response function, there is no noticeable change in the weather response over time.

- Focusing on the morning estimated weather response function, there is a slight lowering in the weather response for both weekdays and weekend days post-2015. The post-2015 weekday and weekend weather response appears to be flatter on the space-heating side and roughly unchanged on the air-conditioning side.
- Focusing on the midday estimated weather response function, there is a lowering in the weather response for both weekdays and weekend days post-2015. The estimated weather response function appears to have steepened on both the space-heating and air-conditioning side post-2015.
- Focusing on the afternoon estimated weather response function, there is a lowering in the non-weather-sensitive portion of the weather response for both weekdays and weekend days post-2015. However, the estimated weather response function appears to have steepened on both the space-heating and air-conditioning side post-2015.
- Focusing on the evening estimated weather response function, there is a slight lowering in the non-weather-sensitive portion of the weather response for both weekdays and weekend days post-2015. However, the estimated air-conditioning response appears to have steepened post-2015 for both weekdays and weekends. The estimated space-eating response appears to be about the same post-2015.

Based on the findings from this analysis it is recommended that the estimated weather response portion of the California ISO production load forecast models be re-specified to capture the electric load evolution from increasing solar PV generation in California. The most straightforward change to the models would be to integrate the weather response portion of the models with a time trend.

CHAPTER 7:
Valuation of Forecast Error Introduction

The goal of this study's valuation effort was to develop and test solar PV forecast economic valuation methodologies. Project partners agreed that the current SolarAnywhere® v4 forecast would serve as the baseline forecast. References to baseline or benchmarks throughout this report refer to this v4 forecast.

This report presents results for three valuation methodologies.

- Resource adequacy
- Real-time market correction
- Perfect forecast

Resource adequacy calculates the cost of capacity based on the standard deviation of forecast errors. Real-time market correction is methodology that values forecast improvements based on prices from the day-ahead market (DAM) and hour-ahead scheduling process (HASP). Perfect forecast is a methodology that values forecast improvements by comparing them with the cost of providing a perfect forecast using energy storage.

The research team studied these three valuation methodologies to approach the problem of valuing forecast errors from three separate perspectives: capacity based (resource adequacy), market based (real-time market correction), and energy based (perfect forecast). In examining the structures of each of these methods, the research team gained insight into the true value of improving forecast accuracy.

Approach

Resource Adequacy

Resource adequacy calculates the cost of reserve capacity required to mitigate some defined percentage of forecast errors based on their empirical probability distribution. The value of moving from one forecast to another is the cost difference between the two.

The California Public Utility Commission's resource adequacy program has two goals, to:

- Ensure that sufficient resources are available within California to provide for the safe and reliable operation of the electric grid.
- Provide targeted incentives for siting and constructing the generation resources needed to ensure future reliability.

The resource adequacy framework ensures that sufficient generation resources are available to meet electric load at all times. Long-term trends like load growth, planned generation, and probabilistic outage modeling are all considered. Part of this framework involves a capacity reserve margin based on load, generation, and transmission characteristics in addition to regulatory requirements. The California Public Utilities Commission ensures that the system remains stable in worst case weather and/or equipment outage events.

The resource adequacy method is based on the probabilistic portion of the resource adequacy framework (the 1-10 weather-year portion). The resource adequacy metric is an estimate of the reserve capacity required to mitigate x percent of forecast errors at a given forecast horizon. The forecast error is the difference between the forecast and what actually occurred. More accurate forecasts yield narrower error distributions and lower reserve capacities and costs.

This metric reflects a more long-term cost assessment than the real-time market correction method (discussed in the next section) because it quantifies the cost of building new capacity.

Forecast errors result from both over-prediction and under-prediction events. Solar PV produces more than predicted during under-prediction events. Solar PV also produces less than predicted during over-prediction events. Under-prediction events are addressed by curtailing solar PV oversupply, ramping down other resources, soaking up excess solar PV with flexible loads, or exporting excess power to other regions. Over-prediction events are addressed by balancing loads with other generation, shedding load, or importing energy from other sources.

This proposed methodology examines the cost impacts of both over-prediction and under-prediction events by identifying the capacity required for ramping up or down to mitigate some percentage of these errors. A capacity cost is then applied so that the cost of errors can be quantified. Calculating the cost difference between the baseline and improved forecasts adds value to this metric. This method yields an upper bound of value because more economic solutions than simply building additional capacity exist to mitigate under-prediction events.

The location of 16 solar PV sites that were included in the study is shown in Figure 33. These locations represent diverse climate regions across California and cover the geographic range of solar PV resources in the state. The metrics described here represent only a small sampling of the installed solar PV in California, and the errors of the actual fleet are assumed to be comparable per unit of generating capacity (per kilowatt of PV capacity). For this reason, the units of this error distribution are expressed in W/m^2 of forecast error per kilowatt (kW) of PV fleet capacity.

Figure 33: Analytic Locations and California Climate Zones

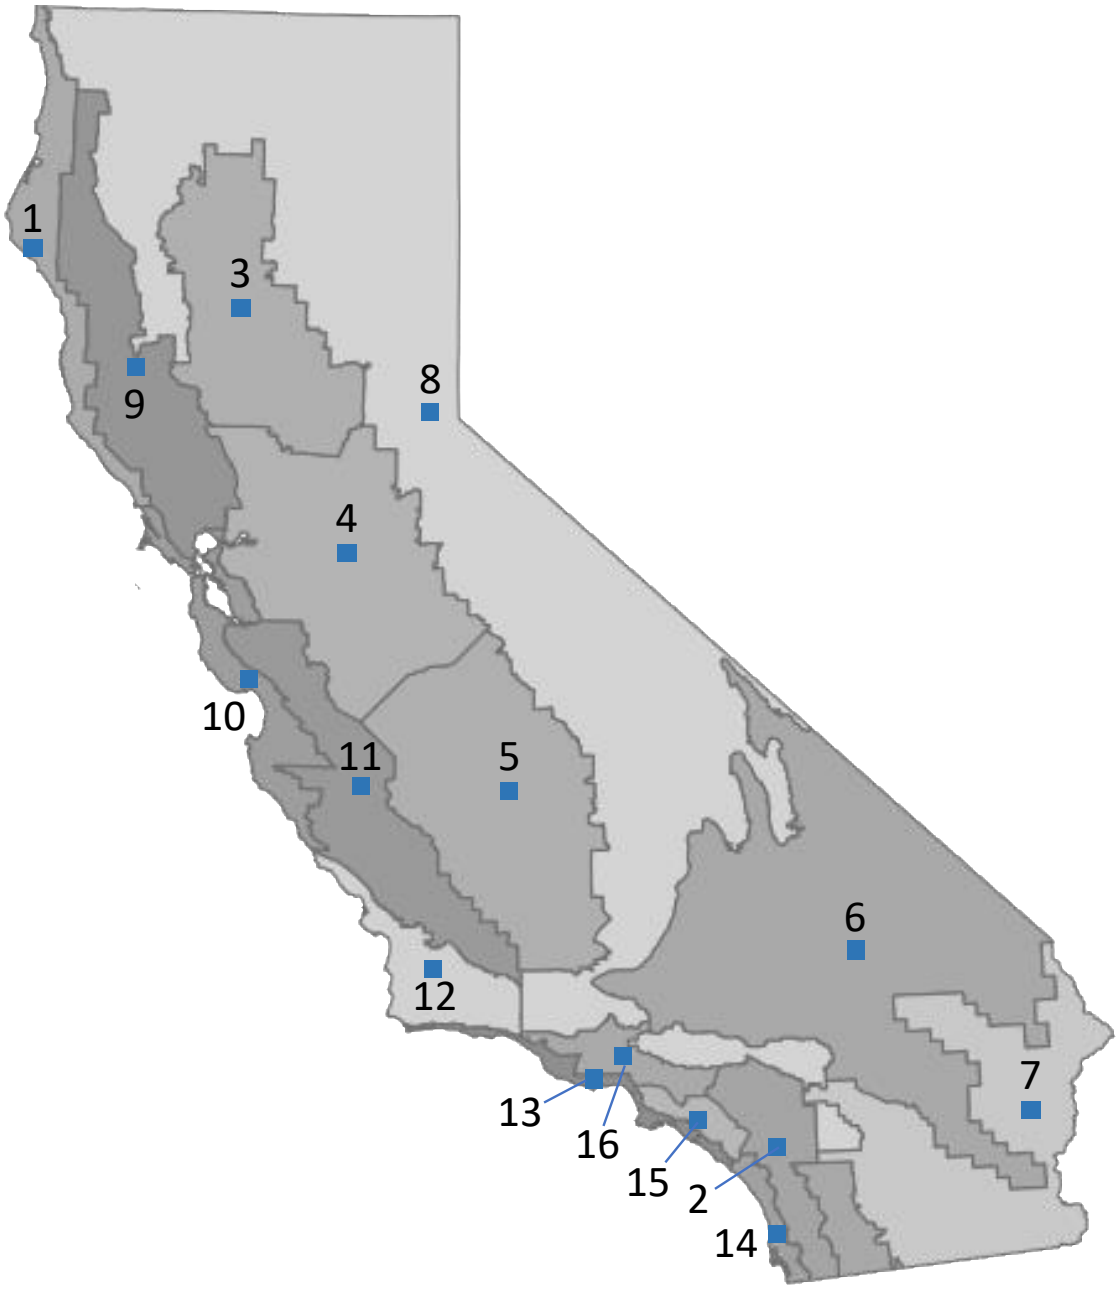

Analytical locations highlighted in blue. Spatial allocation of PV is equally weighted across these blue points.

Source: Clean Power Research

Figure 34 shows the error distribution in the benchmark forecast at 16 locations over the eight-month period data was available. To calculate error, a 1-hour forecast of GHI was made using SolarAnywhere separately for each location, representing the output of a single PV resource. These outputs are added together to arrive at the total fleet production. After the hour passed, the actual fleet output was obtained from the SolarAnywhere historical data set by combining the measured, historical values from all 16 sites. The error represents the error

in GHI forecasts, and does not include error associated with solar modeling, fleet scaling, or other factors outside of the forecasting methods.

Figure 34: 1-Hour-Ahead Forecast Error

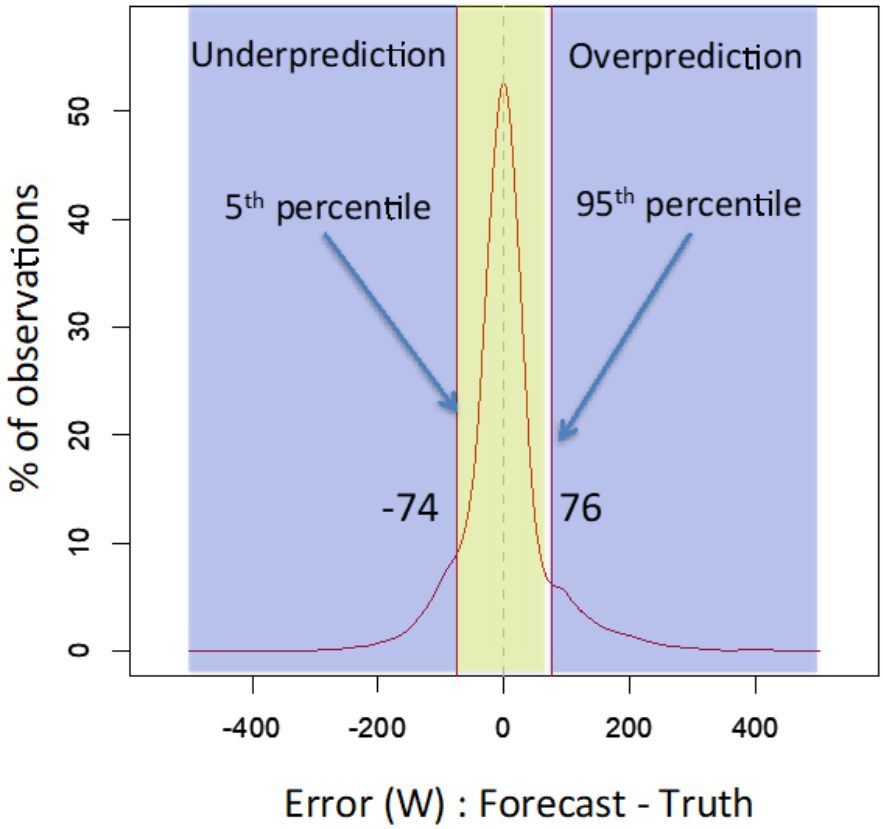

Probability distribution function of hour-ahead forecast errors. Bin width is 50 W/m^2.

Source: Clean Power Research

Irradiance, specifically GHI, is a proxy for solar PV production. This is a reasonable assumption since PV is rated in a sunroom at 1000 W/m^2. This means that 1 kW of PV will produce 1 kW DC of output at 1000 W/m^2. Similarly, 1 kW of PV will produce 500 W DC at 500 W/m^2.

The figure shows the probability distribution of hour-ahead forecast errors across the State of California using the current v4 forecast and a bin width of 50 W/m^2 with empirically-derived 5th and 95th percentile bounds separated by color. Ninety percent of hour-ahead forecast errors fall within the yellow zone and 10 percent fall within the blue zone. This distribution reflects one hour ahead (designated "HA1") GHI forecast errors collected in the hourly interval between July 12, 2015 and March 31, 2016. The width of this yellow zone is 150 W/kW PV. The errors are at a single PV site (point #3 shown in Figure 33).

Real-Time Market Correction

Real-time market correction is a methodology that values forecast improvements based on prices from the California ISO DAM and HASP. The relative value of a given forecast is the cost difference between the baseline (v4) and the improved forecast. Following here is an adaptation of a similar methodology developed by Itron that calculates the value of short-term load forecasts by looking at their cost differences. The original calculation extracts prices from

next-day and day-of wholesale electricity markets and applies them to corresponding day- and hour-ahead forecasts.

The premise behind using short-term market pricing to perform the valuation is that the same amount of electricity is actually consumed regardless of which forecast is used.

The valuation is dependent on three cost components.

- Cost of scheduling resources day-ahead to supply forecasted net load using locational marginal pricing (LMP) from the California ISO DAM
- Cost of making hour-ahead adjustments to the forecasted net load to account for errors in the day-ahead relative to the hour-ahead forecasts using the LMP prices in the HASP
- Cost of making final adjustments to cover the difference between the forecasted load and the actual load using regulation up and regulation down prices from the ancillary services markets, as purchased in the DAM.

Perfect Forecast

The perfect forecast methodology calculates the storage required to mitigate over-prediction events across a given timescale. For example, Figure 35 shows a forecast of what occurred over a single day.

Figure 35: Illustration of Perfect Forecast

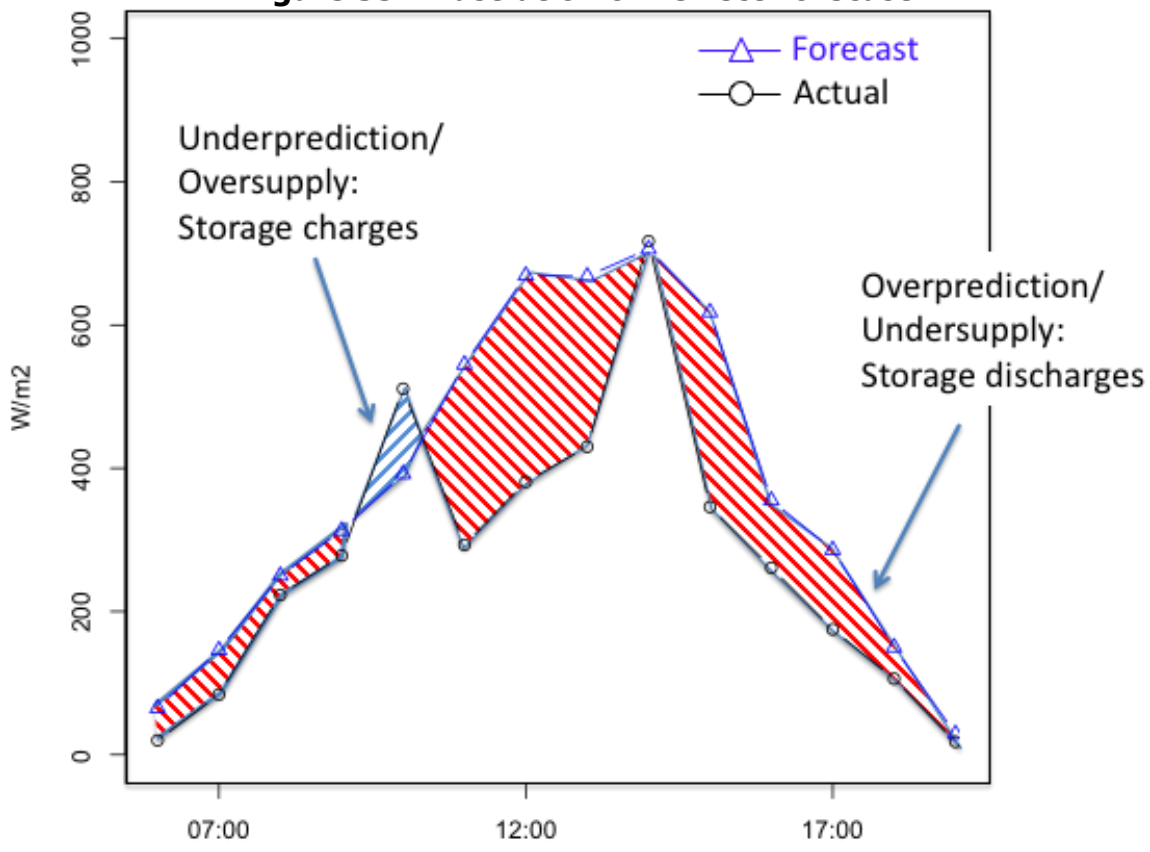

Figure showing a single day's prediction (blue line with triangles) versus what actually occurred (black line with circles). Highlighted are periods of over prediction/undersupply (red hash) and under prediction/oversupply (blue hash).

Source: Clean Power Research

PV is not producing as much as forecasted during an over-prediction event (red areas), so other resources must fill the gap. Currently, this other resource would likely be a quick-ramping fossil-fueled generator. Storage could fill this gap if the goal is to achieve high PV penetration. Storage therefore needs to be sized so that it can accommodate the drawdowns required by these over-prediction events across given time periods.

Results

Resource Adequacy

The research team evaluated four versions of the SolarAnywhere v4 forecast (described in this report's Executive Summary) across four forecast time horizons (Hour Ahead 1 [HA1], HA3, HA24, and HA48) over the time period of July 12, 2015 through March 31, 2016. Figure 36 presents the results.

Figure 36: Resource Adequacy Forecast Error

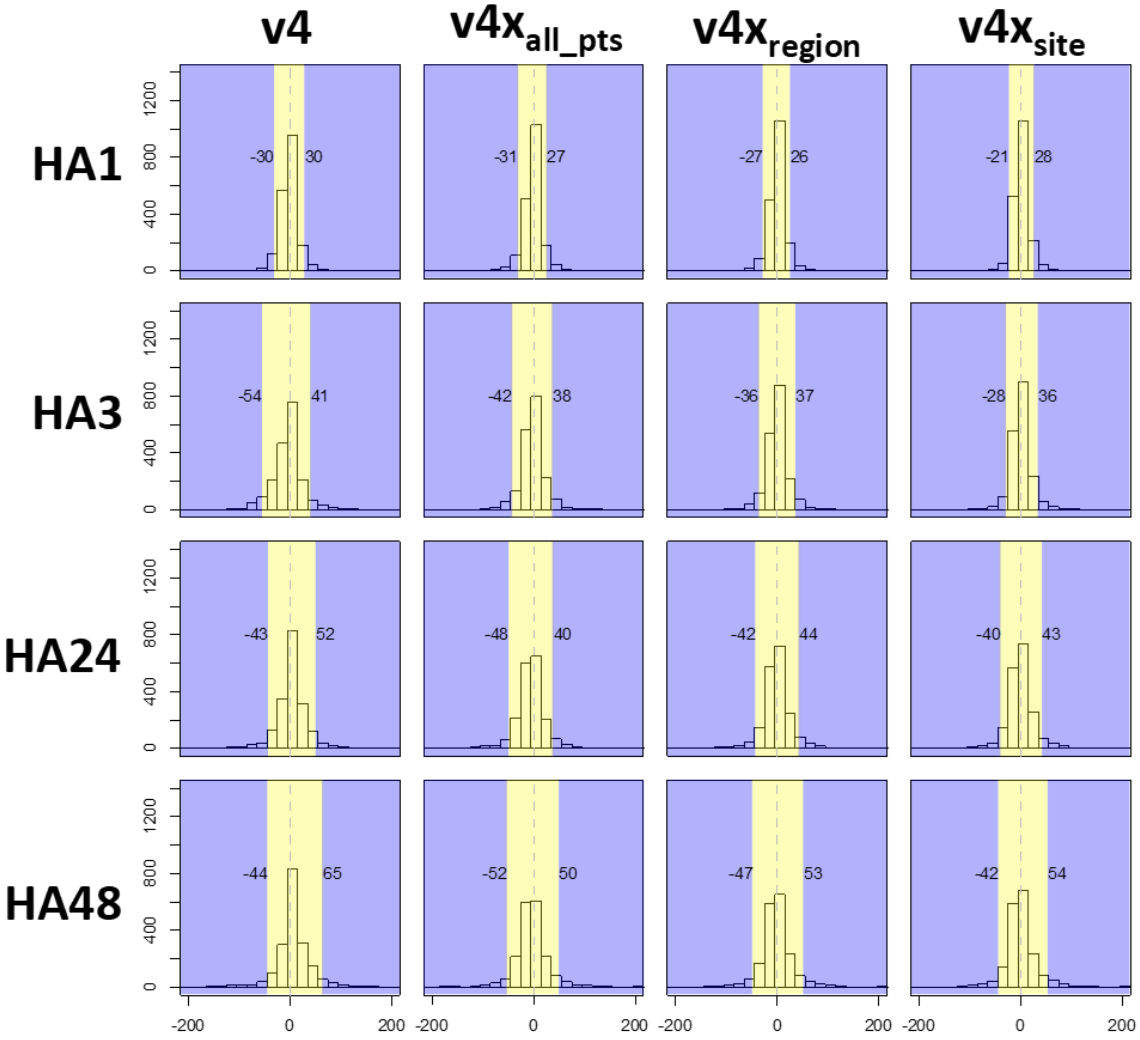

Histograms of the forecast error (forecast – actual in W/kW PV) for each of the four forecast horizons and across each of the four forecast products provided by SUNY

Source: Clean Power Research

Each histogram has a yellow-shaded region and corresponding text that includes 90 percent of the results and errors. All the results shown here reflect an even spatial distribution of PV across each of California's climate zones. The left column of plots represents the distributions for v4 forecasts, while the next columns represent v4x with all-points weighting, regional weighting, and site weighting, respectively.

Figure 37 summarizes results from the 16 plots in Figure 37, using two vertical axes. The left axis corresponds to the 90-percent bound in W/kW PV. The right axis corresponds to the cost of this capacity in $/kW PV, given a capacity cost of $0.88/W. Each forecast horizon has its own color, and each cluster of columns represents a different forecast version.

The reserve capacity and corresponding cost per unit of PV:

- Decrease with increased geographic site specificity (bars are smaller for v4x_site than they are for v4); and
- Increase with lengthening forecast time horizons.

For every forecast horizon, each of the v4x forecasts shows improvements in the load forecast under consideration.

Figure 37: Resource Adequacy

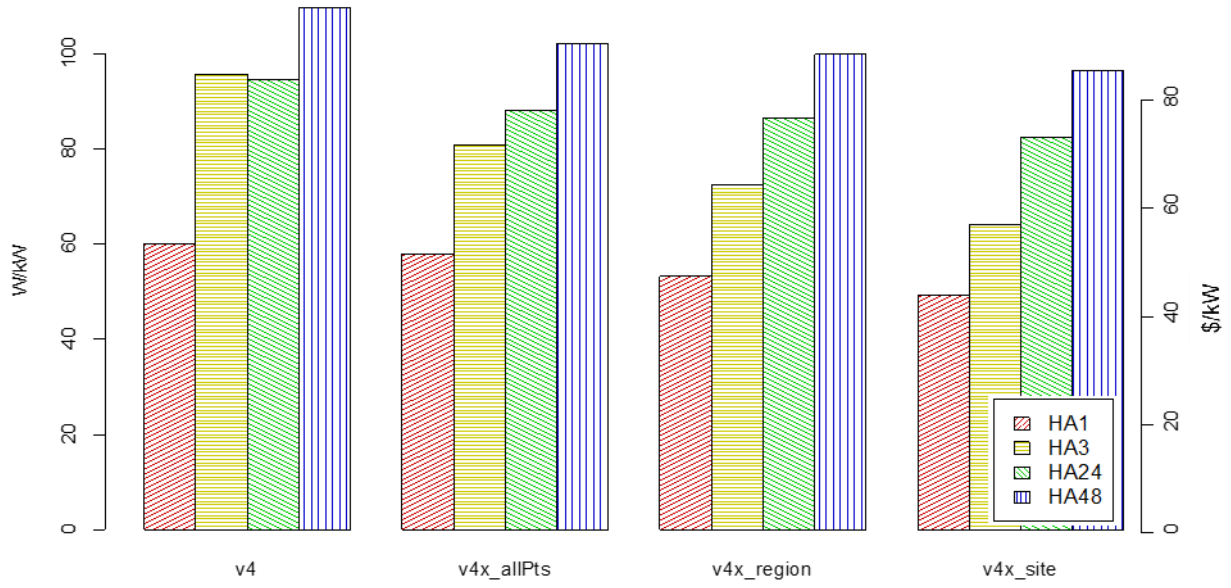

Resource adequacy in W/kW (left axis) and $/kW (right axis) based on the width of the 90-percent bounds for each of the four versions of the forecasts and at each of the four forecast horizons. Capacity costs of $880/kW reflect a modern-day advanced combined-cycle gas generator.

Source: Clean Power Research

Figure 38 shows the savings when moving from v4 to each of the v4x versions of the forecast in $/kW. The savings increase with greater geographic specificity of the forecast product. The capacities presented in Figure 37 and corresponding costs in Figure 38 reflect an even geographic distribution across all 16 points in California. These values are roughly half of what is illustrated in the methodology section, which reflects a single one of these locations. This halving of costs reflects the effect of geographic dispersion.

Figure 38: Resource Adequacy Savings as Compared to v4 Forecast

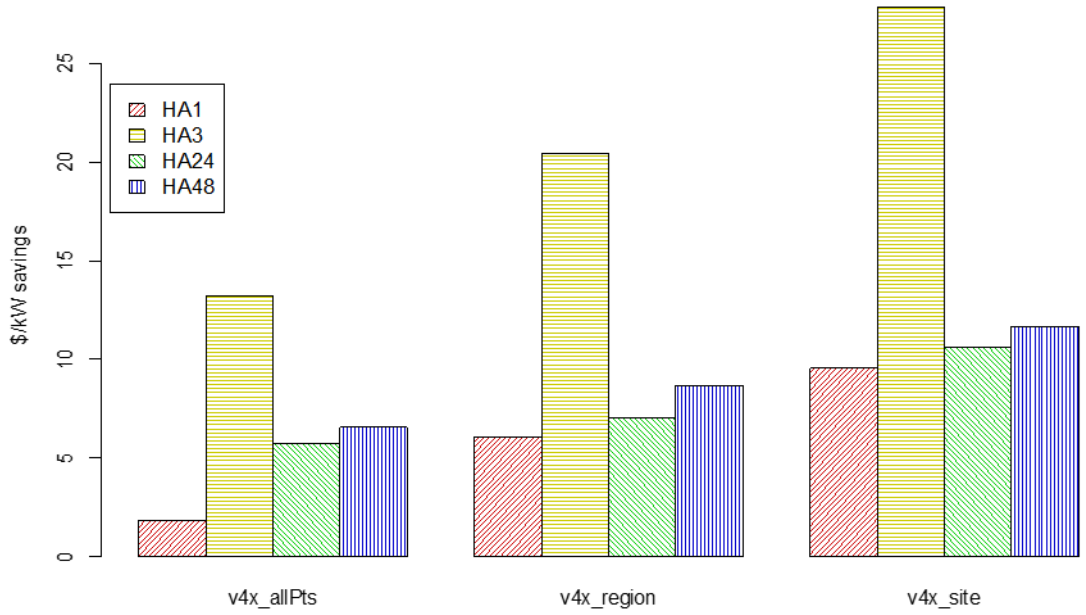

$/kW savings when employing each of the three improved versions of the forecast relative to v4, the forecast currently supplied to the California ISO.

Source: Clean Power Research

The savings offered by the version v4x_site over the baseline v4, for example, are $28 per kW of PV. Translating the shift from v4 to any of the three v4x versions at various time horizons into dollars for a given capacity of PV was performed. The expected savings to California when shifting the HA3 forecast version from v4 to v4xsite is $176 million (31.8 W Natural Gas Combined Cycle / kW_PV x 0.88 $/W Natural Gas Combined Cycle x 6,300,000 kW PV in CA) in terms of avoided combined cycle gas turbine (CCGT) backup capacity. This is the amount of money required to build new gas capacity capable of erasing 90 percent of 3HA forecast errors for the state's PV fleet.

This example is predicated on uniform cost scaling of forecast error from the 16 sites analyzed in the encompassing region. The assumption here is that the forecast errors within each region entirely correlate with one another. Therefore, the aggregate error when combining the entire BTM PV fleet, for instance, will mimic the aggregate error across the 16 sites. This isn't entirely true; in reality, some of the sites within each individual region will correlate less since many microclimates exist in California.

Real-Time Market Correction

To model relative forecast value by the methodology outlined above, it was necessary to obtain time-synchronous price data from the California ISO's open access same-time information system (OASIS). The research team used the application program interface as outlined in version 5.0.0 of the OASIS interface specification to obtain hourly price data from July 11, 2015 through March 31, 2016. This temporal range corresponds to the sample used for each of the three valuation methods.

Figure 39 and Figure 40 show price components from OASIS across the period considered. The LMP's shown represent only the energy component within both the DAM and HASP. This energy component is the same across all nodes, simplifying the calculation.

Figure 39: Open Access Same-Time Information System Prices

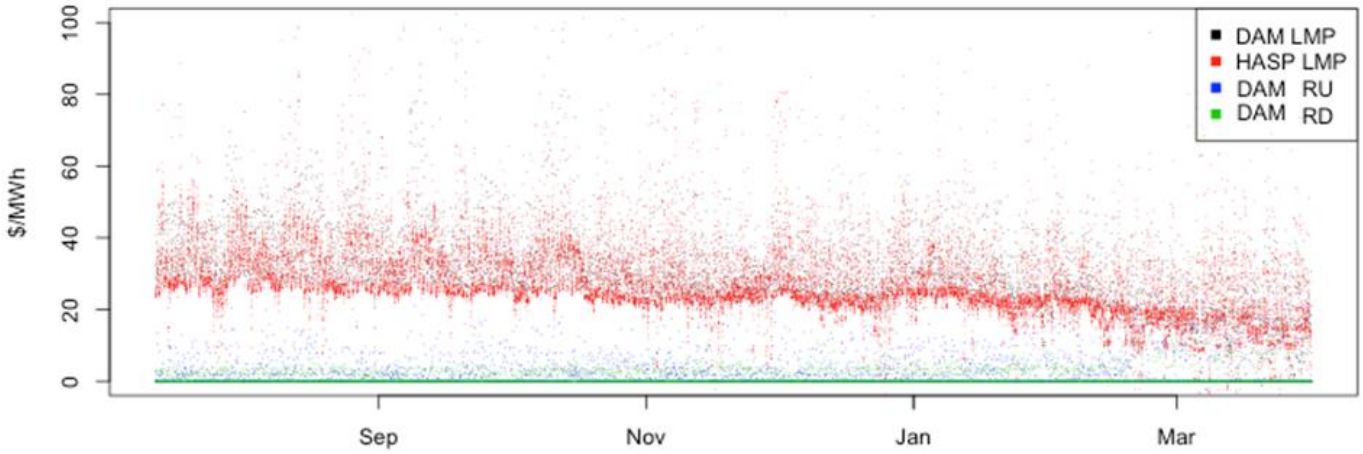

Prices from OASIS over the period of interest. Though the y-axis is cropped to 100 $/MWh, the range of the HASP LMP extents upwards to 1000 $/MWh across this period. DAM RU refers to day-ahead market, regulation up; and DAM RD refers to day-ahead market, regulation-down.

Source: Clean Power Research

Figure 40: Open Access Same-Time Information System Prices for Sample Week

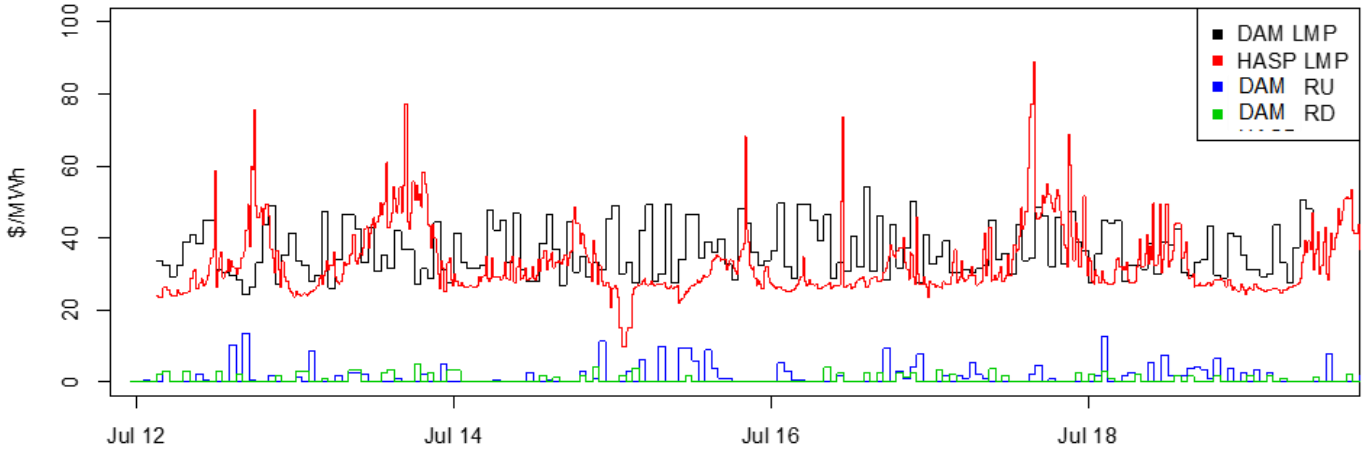

Zoomed-in series of prices from OASIS across a single week. Note how the temporal frequency of the HASP market (15 minutes) is visibly more granular than the frequency of the elements pulled from the DAM (1-hour).

Source: Clean Power Research

Figure 41 presents two x-y plots of data pertinent to calculating forecast value: 1-hour-ahead forecasts (HA1) and 24-hour-ahead PV forecasts (HA24) versus historical data across all of California. As expected, the HA24 forecast is less accurate than the HA1 forecast.

Figure 41: Measured Versus Forecasted Global Horizontal Irradiance

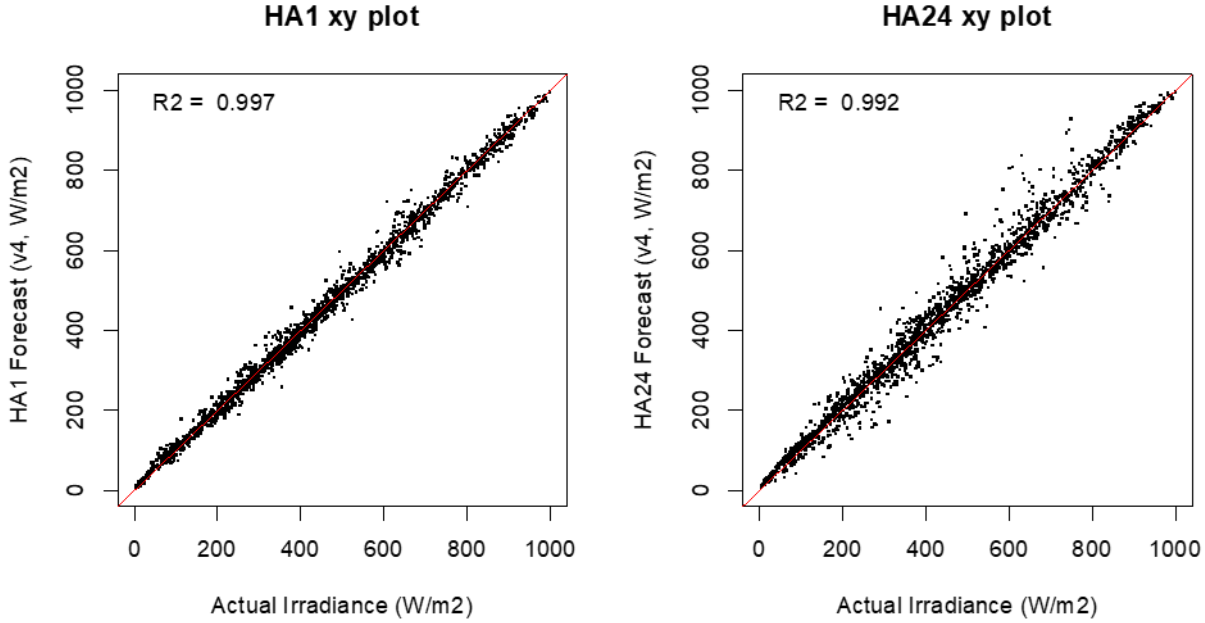

Global horizontal irradiance (actual) in W/m2 versus HA1 and HA24 forecasts, respectively. Correlation values are listed in the upper left-hand corner.

Source: Clean Power Research

All of these data were pulled together and timestamps were aligned. The granularity of the HASP market data was at the 15-minute interval. It was therefore extended to match the hour interval of the DAM data by time averaging. The research team tested the methodology by modeling the 6.3 gigawatt (GW) of BTM PV installed in California as a fleet evenly distributed across the state. Figure 42 presents the result of this analysis. Savings are in millions of dollars per year and the corresponding values are shown in Table 6. The highest value for an even spatial distribution of 6.3 GW of PV is $0.3M/yr if the research team switches to v4xsite from v4.

Table 6: Value of Forecasts

v4xallPts	v4xregion	v4xsite
-1.2	-0.3	+0.3

Value in $M of switching to one of three alternative forecasts from v4 (baseline) using this methodology

Source: Clean Power Research

The improved forecast version v4xsite weighting yields the highest value using this metric if PV is spread evenly across the state. Switching the HA1 and HA24 forecasts to v4xsite yields a value of roughly $300k/yr. Interestingly, a negative value results when switching to v4xall.pts or v4xregional. This is because the DA forecast in v4 is larger than the DA forecast for v4xall.pts or v4xregional, making the first term of the value equation weigh heavily.

Figure 42: Forecast Value

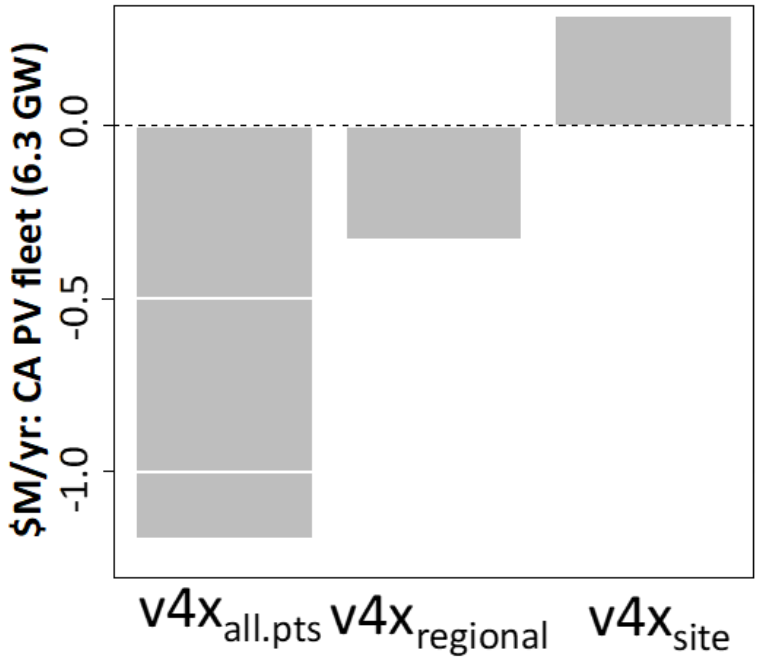

Column plot showing net value of switching from current in-operation v4 forecast to one of three versions of the improved forecast if PV is geographically allocated evenly across California. Values are in $M of dollars per year when considering a PV fleet size of 6.3 GW.

Source: Clean Power Research

Perfect Forecast

Assume a future storage cost of $100/kWh capacity. Figure 43 presents the savings from moving from the v4 forecast to the three improved forecasts. It is based on forecast data for a single point. The figure shows that more money is saved as the geographic specificity of the forecast tuning is increased (moving from left to right on the graphic plot).

Figure 43: Forecast Improvements

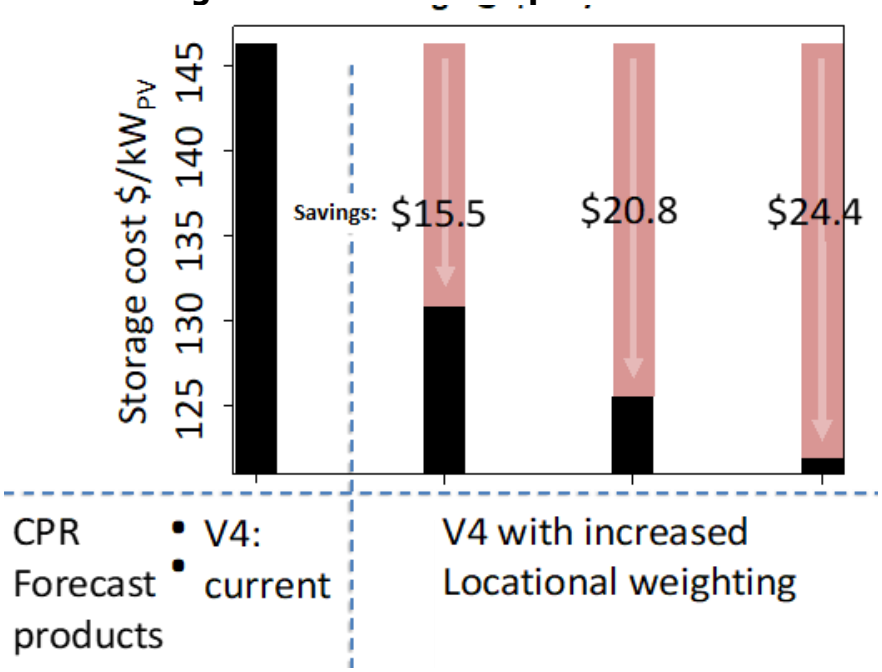

Nominal storage costs in $/kW PV for the currently-in-operation v4 forecast and all three of the improved forecasts studied in this project. Forecast value relative to v4 is highlighted above each bar. Storage costs are assumed to be $100 per kWh.

Source: Clean Power Research

Note that the research team calculates cents per kWh instead of nominal dollars per kW of capacity in the other sections. This is done because with this method alone the research team can put a price on curtailed electricity from under-prediction and oversupply events.

Table 7 presents costs from a cents-per-kWh perspective. What is the cost of guaranteeing various versions of the forecast in c/kWh? For these calculations, the research team again used a storage cost of $100/kWh and a PV cost of $300/kW. Given these assumptions, meeting all HA1 forecast errors with PV + storage will cost 2.63 c/kWh. This is 15 percent more expensive than the cost of PV alone (2.28 c/kWh) because of the storage required to perfect the forecasts. Switching to v4xall.pts saves 0.01 c/kWh, while switching to v4xregion or v4xsite increases costs slightly for the HA1 forecast. Switching to any of these three forecast types for the HA3, HA24 or HA48 horizon saves money relative to the v4 forecast.

Table 7: Cost of Energy and Associated Premiums to Back Up with Storage

	HA1				HA3				HA24				HA48			
	v4	v4x			v4	v4x			v4	v4x			v4	v4x		
		all.pts	region	site		all.pts	region	site		all.pts	region	site		all.pts	region	site
COE (c/kWh)	2.63	2.62	2.68	2.84	3.4	3.08	3.11	3.34	5.52	3.7	4.22	4.15	6.42	3.94	4.2	4.27
Premium (%)	15	15	18	25	49	35	37	47	142	62	85	82	182	73	84	88

Cost of energy in cents per kWh using future capital costs for PV and storage, and premium over COE of PV without storage

Source: Clean Power Research

Conclusions and Recommendations

Resource Adequacy

If new gas generation were built expressly to accommodate 90 percent of the HA3 forecast errors across California's 6.3 GW BTM PV fleet, this metric indicates savings of up to roughly $176 million. These savings accrue when moving from v4 to the newest, most geographically-tailored version (v4xsite).

The resource adequacy method has two limitations. First, it gives equal weight to under- and over-prediction errors when forward and reverse bias errors have different cost profiles. Second, the capacity required to mitigate forecast errors will have a low capacity factor if it is used solely to mitigate PV forecast errors. If this reserve capacity were actually built, it would likely play other grid-service roles as well. It is therefore an overestimate to attribute all of the $0.88W cost to forecast-error mitigation alone.

Real-Time Market Correction

Using an adapted version of the methodology defined by Itron, the research team calculated that changing from the v4 forecast to the v4xsite forecast will save $300,000 per year across the state.

This metric is probably the most complex of the methods described in this report because it requires time-correlated market data. Because of this use of market data, it is perhaps more intrinsically palatable to grid operators and other utility stakeholders who are already familiar with it.

One issue with this metric is that forecasts can be manipulated to intentionally show high value.

Note that in Figure 42 and in the corresponding Table 6, two of the three new forecasts (v4xall.pts and v4xregional) examined present negative values when using this method. This does not mean that these are unimproved forecasts (they are by measures of RMSE, forecast skill and the two other valuation methods discussed in Chapter 1 and Chapter 3. The "game-ability" of this method indicates that this method is not useful for calculating value.

A final insight is in relation to costs linked to forecast scaling. The most valuable forecast using this metric is the one biased for underprediction. In this case, additional resources are brought online to mitigate the shortfall, and this is less expensive that forecasts biased for overprediction because ramping down is less expensive than ramping up.

Perfect Forecast

The perfect forecast metric reflects the fundamental energy balance between what was predicted and what actually occurred.

The metric quantifies the amount of storage required to guarantee forecasts at a specific time horizon. It sets a high bar because, as a firm-power guarantee, it penalizes the most extreme cumulative over-prediction events. Unlike the resource adequacy metric or the real-time market correction metric, perfect forecasts force decisions not only in terms of instantaneous power, but also in terms of energy across extended periods. This is important because renewable energy is plentiful but doesn't share the same ramp-up, on-demand characteristics

of conventional generation. This method quantifies that incremental cost if storage is required to guarantee PV forecasts.

By scaling forecasts, the research team reduced the amount of storage needed to mitigate forecast errors regardless of the horizon. By paying just a 9.3-percent premium per kWh for curtailment and storage, the research team mitigated all hour-ahead forecast errors. This is a reduction from the 15-percent premium the research team would theoretically pay if no curtailment optimization were applied to correct hour-ahead errors shown in Table 7. The research team mitigated all 48-hour-ahead forecast errors with storage by paying a 28.5-percent premium per kWh.

CHAPTER 8:
Distributed Energy Resource Capacity Forecasting

Introduction

As solar PV and other DER technologies become available and more affordable, retail electric customers are increasingly turning to these on-site energy alternatives.

Both residential and commercial utility customers are installing rooftop solar PV in increasing numbers. Energy storage options are also increasingly available and affordable for on-site energy management, propelled in part by high manufacturing volume in the EV sector. EVs themselves, with high charging consumption and controllable charge rates, are, in aggregate, affecting grid management. New options for electrification such as electric heat pumps for space heating and water heating, are also becoming available, replacing natural gas furnaces and water heaters. Smart thermostats and other smart devices now provide the means to shift flexible thermal loads to lower-cost time periods.

Customer decisions to adopt DER technologies are motivated by a variety of factors. Some may seek to lower their utility bills or take advantage of federal tax incentives and short payback periods. Some may seek clean-energy alternatives to conventional fossil fuels due to personal environmental beliefs or a desire for energy independence. Some may be attracted to new technology, and some may respond to a combination of any of these factors.

Whatever their reasons, the large number of utility customers adopting these DER technologies are significantly affecting California's electric utilities. Overall grid-electricity usage decreases with on-site solar PV generation but increases with EV charging. Both impact utility revenues and, consequently, electric rate design. Depending on the technology, the timing of daytime electricity usage can shift. Battery dispatch, smart-thermal load control, and EV-charging control all modify demand timing, ultimately affecting resource planning and changing wholesale electricity market operations. DER electrification decreases gas usage but increases electricity usage, which affects the financial outlook for both gas and electric utilities.

It is not surprising then, that long-term DER forecasting methods play an important role in the state's energy future. Accurate forecasts of solar PV capacity coming online will allow utility resource planners to better gauge the amount and types of new generation required to meet customer demand. Forecasts of new EV load and its charging demands will help electric resource planners and distribution engineers plan for additional distribution capacity and system protections. Hourly models of DER forecasts will enable rate designers to accurately design demand charges, energy charges, and time-of-use periods. Financial models, enhanced by forecasted DER adoption, will provide utility analysts with the data they need to estimate financial impacts to utility shareholders, customer-owners, and municipal-utility customers. DER capacity forecasts will enable energy policymakers to identify risks, design incentive programs, and plan for clean energy advances.

Approach

CPR developed a logistic regression model using customer-level input data to forecast customer adoption of DER technologies. In this model, input data can be defined separately for each technology type. For example, to forecast PV adoption, inputs could include customer payback periods, demographic segments, or local customer adoption densities to date. The battery energy storage adoption forecast may include whether the customer has previously adopted PV and the number of prior outages on a local circuit. Electric vehicle adoption may include a binary input describing whether the customer evaluated the EV on a utility-provided informational tool. Other inputs could be considered based on data availability and tests of their significance, described in the next section.

Model Overview

The input data was first combined in a linear model, calculated for each customer by combining each input predictor variable (explanatory variable), weighted by a coefficient and a stand-alone intercept term. In logistic regression, the linear model represents the logarithm of the odds of adoption, or "logit," thus defining the link function from which the probability is calculated.

Model coefficients in logistic regression are determined using any of several statistical packages, such as those available in R or Python. These tools select coefficients that best fit customer responses to available training data. This is accomplished with the likelihood function, illustrated in Figure 44. The figure shows customers in the training data set as circles, with g(X) indicating the linear model, and with response coding "1" for adopters and "0" for non-adopters.

For a given set of coefficients and intercepts, the transformed model (red line) is the probability of adoption $\pi(X)$ as a function of g(X). The likelihood function is built from all customers; for each adopter, such as customer "A," a factor $\pi(X)$ is assigned, which is the probability of adoption. For each non-adopter, such as customer "B," a factor $1-\pi(X)$ is assigned, which is the probability of not adopting. The likelihood function is the product of all factors across the customer set. The likelihood function is therefore a measure of fit, so the coefficients and intercepts maximize the likelihood function.

Figure 44: Building the Likelihood Function

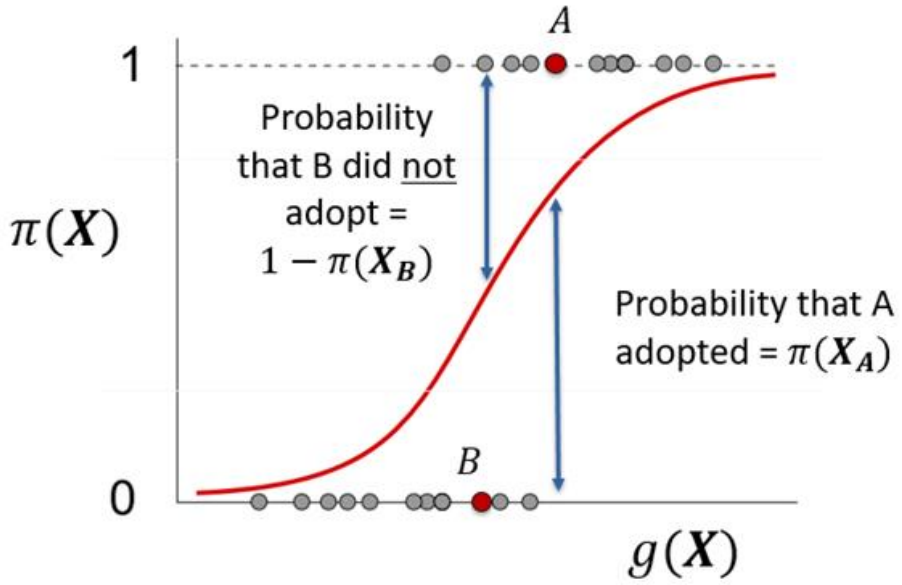

Source: Clean Power Research

Future-year probabilities were calculated using the same model coefficients and suitable adjustments to the predictors themselves. For example, in year 10, input predictors could be modified using scenario assumptions about future installed-DER costs, electric rate structures, availability of federal tax credits, and other factors. Adoption density in year 10 could then be calculated from existing installed DERs and the cumulative probabilities of adoption for all customers.

Customer Data Set

Input data for approximately 600,000 electric customers in Northern California was used to validate the PV forecast model. Data included hourly loads for 2016, rate class, PV commissioning date, market segment, and location. From this, CPR developed and tested a logistic regression PV adoption forecast model for residential customers.

Predictor Variables

Predictor variables were defined for cost-effectiveness, adoption density, and market segment. The cost-effectiveness variable was the inverse payback; payback was calculated using year-specific historical or future-year estimated PV installed costs, a calculation of annual bill amounts from hourly loads and time-correlated PV production, and year-appropriate incentives.

Adoption density was defined as the prior-year cumulative PV adopter count in a 10-mile x 10-square-mile region, centered at customer location. This is illustrated in Figure 45. If the model training year was t, then all prior adopters in year t-1 were counted in 2-mile x 2-mile cells throughout the service territory. In the illustration, three such adopters were found in the cell containing the customer, and 76 such adopters were found in the cells in the 5x5 matrix centered on that cell. The predictor variable for that customer in year t would therefore be set at 76.

Figure 45: Illustration of Adoption Density Calculation

4	0	2	1	2	3	0	1	1
2	0	0	3	5	1	5	0	2
3	2	0	1	3	0	0	1	1
2	4	2	3	0	6	2	8	4
4	0	3	0	3	3	0	2	4
1	4	3	3	8	4	8	5	0
0	5	3	3	2	9	7	6	3
4	4	3	0	3	10	9	6	1
3	2	2	3	8	9	7	6	2

Source: Clean Power Research

The market segment predictor was defined as a continuous score based on the number of prior PV adopters in the customer's segment. For example, there were 10,512 customers in one specific segment. Of these, 151 adopted PV (1.43%) by the prior year. Percentages among all segments ranged from 0% to 3.37%. So, customers in this segment were assigned a predictor variable value of 0.0143/0.0337 = 0.424.

Training

CPR used customer data for 2014 and made a forecast for 2015 adoption. Having actual adoption data for both years enabled accurate forecast calculations. Only residential customers were included except in calculations of adoption density since, by definition, both residential and commercial PV systems were used in adoption density calculations.

Results

When a customer installs a DER technology, it changes that customer's hourly load profile. Production patterns from new PV capacity follow changes in local meteorological conditions, so charge and discharge storage rates depend upon current states of charge and storage dispatch algorithms. Each DER technology similarly has a profile, potentially including both the production and absorption of energy. This energy represents net load.

Modeling DERs also depends upon both physical sizing constraints and design attributes. In the case of PV, modeling depends upon the kW power rating of the system and other parameters including the tilt and azimuth angles of the panels. In the case of storage, maximum kW charge and discharge ratings must be determined, as well as energy storage capacity (kWh). Each technology has similar required attributes to model its output.

Hourly production modeling, such as the California IOUs' PV fleet production profiles, is subject to three broad categories of potential error. First, there are errors associated with forecasts of customer adoption. This error may follow from general difficulties in modeling complicated consumer behaviors, lack of input data sources, or errors in input data sources.

Second, there may be errors associated with the technology's assumed design attributes. A 20-kW PV system will have twice the hourly output of a similarly configured PV system rated

for 10 kW. So, if the assumed average system forecast size is 20 kW when it is actually 10 kW, this would lead to an error of 100%. Similar modeling errors for all technologies will persist in system design attributes. For example, there will be errors introduced by assuming that a system is oriented due south while, in fact, it is oriented in a south-east direction. Similar errors would be associated with other design attributes, such as tilt, inverter efficiency, and loss factors.

Finally, there are errors in production modeling itself. Inaccuracies in local irradiance data will cause PV production modeling errors. The assumed battery or EV charging algorithm will also result in estimated charging or discharging patterns that may differ from current patterns.

Only errors related to the first type — DER adoption forecast errors — are covered in this section. This simplifies the analysis since the errors discussed follow only DER adoption modeling that is the subject of this report. Errors of the second and third types — design assumptions and production modeling — would be quantifiable using the methods described in the "Valuation of Error" chapter. Such an analysis could be performed, theoretically, a year after the installations. Hourly measurements could then be taken on these future installed systems and the actual cost of forecast error could be quantified.

Test of Significance

CPR developed defined variables based on consideration of what data could be good predictors of adoption. In general, variables that correlate or anti-correlate with adoption rates should be good predictors. However, the question remains as to whether a defined input variable actually improves the model. If a variable does not improve the model, then there is no benefit in including it.

The coefficients themselves do not indicate significance. They are rather "calibration constants" that depend on predictor units and their definitions. Statistical libraries used for logistic regression include built-in inference tests for significance. CPR used the Wald test in this analysis.

Here's how the test works. The "null hypothesis" is when the predictor variable cannot indicate adoption. The research team therefore compared a model that includes the variable in question with a model where the variable was excluded. The comparison was therefore made using likelihood.

The Wald test calculates the ratio of maximum likelihood of the variable's coefficient to an estimate of its standard error. This follows a normal distribution. If the probability was small, the research team rejected the null hypothesis, and the predictor was said to be statistically significant. For example, if a two-tailed p-value for a given predictor is $Pr(>|z|) < 0.001$, it means that there is less than a 0.1 percent chance that the null hypothesis is correct. The research team could therefore conclude with high confidence that its predictor is statistically significant.

While arbitrary, it is common practice to include all parameters with p-values less than 0.05. This practice was followed here. Results from the significance tests follow.

Goodness-of-Fit

Even variables shown to be highly significant do not necessarily guarantee that the model closely fits the observed data. The test of significance only indicates that including the variable will result in a better model than excluding that variable. A test is therefore required to determine how good the model is overall in fitting modeled results to actual results when all predictor variables are included.

In linear regression, R^2 is a commonly used metric for goodness-of-fit. In logistic regression, there are several analogous alternatives. CPR selected the following formula:

$$\text{Goodness of Fit} = 1 - (\text{Residual Deviance})/(\text{Null Deviance})$$

This formula is analogous to R^2 since the value is always between zero and one, and the closer it is to one, the better the fit.

Null deviance measures how well the observed responses compare with the simplest of models with no predictors and only an intercept. In such a model, all customers are equally likely to adopt. Residual deviance is a measure of how well observed responses compare with actual responses, using all predictors and their coefficients.

Aggregate Forecast Accuracy

Finally, CPR defined a simple measure of forecast accuracy as the error in the aggregate number of forecasted adopters. If a forecast estimated the number of adopters to be 105 in a given year and the actual number was only 100, this measure indicates an aggregate forecast error of 5 percent.

Adoption Forecast

Results are shown in Table 8. The first two columns show predictor term names and symbols. The third column shows the corresponding coefficients that maximize the likelihood function. The last columns show the results of the significance test, the goodness-of-fit test, and the aggregate forecast error.

All three predictors are shown to be statistically significant. Asterisk codes are shown only to reflect general significance levels. All predictors are statistically significant since they meet the criteria that p-scores must be under 5 percent.

The goodness-of-fit metric is calculated as 0.056. This shows that the residual deviance (20,212) is better than, but not much better than, the null deviance (21,404). This suggests that it is difficult to determine high differentiation probabilities between customers.

The forecast error is -1.8%, strong confirmation that the model is very good at predicting overall adoption for the residential customer base.

Table 8: PV Adoption Forecast Results

	Predictor	Coefficient	Significance Pr > |z|	Null Deviance Resid. Deviance Goodness-of-fit	Forecast Error
	Intercept	-10.19	< 2e-16 ***	Dn = 21404 Dr = 20212 1 – Dr/Dn = 0.056	-1.8%
PV Cost-effectiveness	X_1	36.78	< 2e-16 ***		
PV Adoption Density	X_2	-3.92e-04	0.00107 **		
PV Market Segment	X_3	2.573	< 2e-16 ***		

Source: Clean Power Research

Conclusions and Recommendations

CPR developed a logistic regression model that forecasts DER technology adoption and demonstrates its PV adoption. Predictor variables were defined to reflect customer cost-effectiveness, local adoption density, and market segment. All variables are statistically significant, but the ability of the model to differentiate customers by probability of adoption is not high. Overall, the forecast agreed within 2 percent of the actual adoption rate.

This model provides the versatility required to forecast any DER technology, provided that suitable input data is available and adoption rates can be measured for training purposes. The model could be used to forecast batteries, EVs, flexible loads, and other technologies.

Methods are described for calibrating the model, using best practices for logistic regression, applying the resulting model coefficients to calculate probabilities, testing for candidate predictor variable significance using the Wald test, measuring goodness-of-fit based on null deviance and residual deviance, and measuring aggregate forecast accuracy.

CHAPTER 9: Technology/Knowledge/Market Transfer Activities

Overview

To transfer the technology and knowledge to the energy industry, CPR and its partners (Itron, UCSD, Jim Blatchford, and SUNY) presented this project's results to various industry conferences, published a journal article, provided data to a California DG Statistics website, and sent the project task report concerning California ISO integration to energy system operators both in the US and abroad.

Conferences

A summary of conference presentations follows. These conferences were all industry-sponsored events, so specialized professionals attended. The Energy Systems Integration group meetings included energy system operators specifically engaged in renewable resource forecasting and grid integration.

- American Meteorological Society Meeting, January 7-11, 2018, Austin, Texas: Wu E, Clemesha RE, Kleissl J. "Coastal Stratocumulus Cloud Edge Forecasts"
- World Conference on Photovoltaic Energy Conversion (WCPEC-7), IEEE Photovoltaic Specialist Conference, June 10-15, 2018, Waikoloa, Hawaii: R. Perez, et.al, "A New Version of the SUNY Solar Forecast Model: A Scalable Approach to Site-Specific Model Training" (presented by M. Perez)
- Energy Systems Integration Group Meeting, June 19, 2018, Minneapolis, Minnesota: T. Hoff, "Extending Fleet Forecasting Capability into the Probabilistic Realm"
- American Solar Energy Society, Solar 2018 Conference, August 5-8, 2018, Boulder, Colorado: R. Perez et al., "PV Fleet forecasting — A New Version of the SUNY Model With Localized Satellite Training Scalability."
- Association of Edison Illuminating Companies, March 26-29, 2019, Orlando, Florida: B. Norris, "DER Forecasting Methods"
- 13th Annual ISO/RTO/TSO Summit, Itron User Group, May 6-7, 2019 Denver, Colorado: J. Blatchford, "Improve Irradiance Forecast Accuracy," (Task 2); F. Monforte, "Integrate Holistic Forecast into California ISO Operations," (Task 5)
- Energy Systems Integration Group Meeting, June 4-6, 2019, Denver, Colorado: M. Perez, "Quantifying Forecast Error"

Journal Article

A peer-reviewed journal article was published by UCSD in the journal *Solar Energy*. This journal is the official journal of the International Solar Energy Society®, which is exclusively devoted to the science and technology of solar energy applications.

The article is available at:

https://www.sciencedirect.com/science/article/pii/S0038092X18302056

CaliforniaDGStats

Project data developed for the DER Production Database were made publicly available on the California Distributed Generation Statistics (www.californiadgstats.ca.gov) website. This dataset updates and corrects an existing set of measured PV system production data from 411 BTM systems throughout California. The data set was produced using specifications from the inference engine, CPR's historical location-specific solar resource data, and CPR's PV simulation model. Supplemental material, including this report, is also available for download. A screenshot is shown in Figure 46.

Figure 46: Screenshot from CaliforniaDGStats

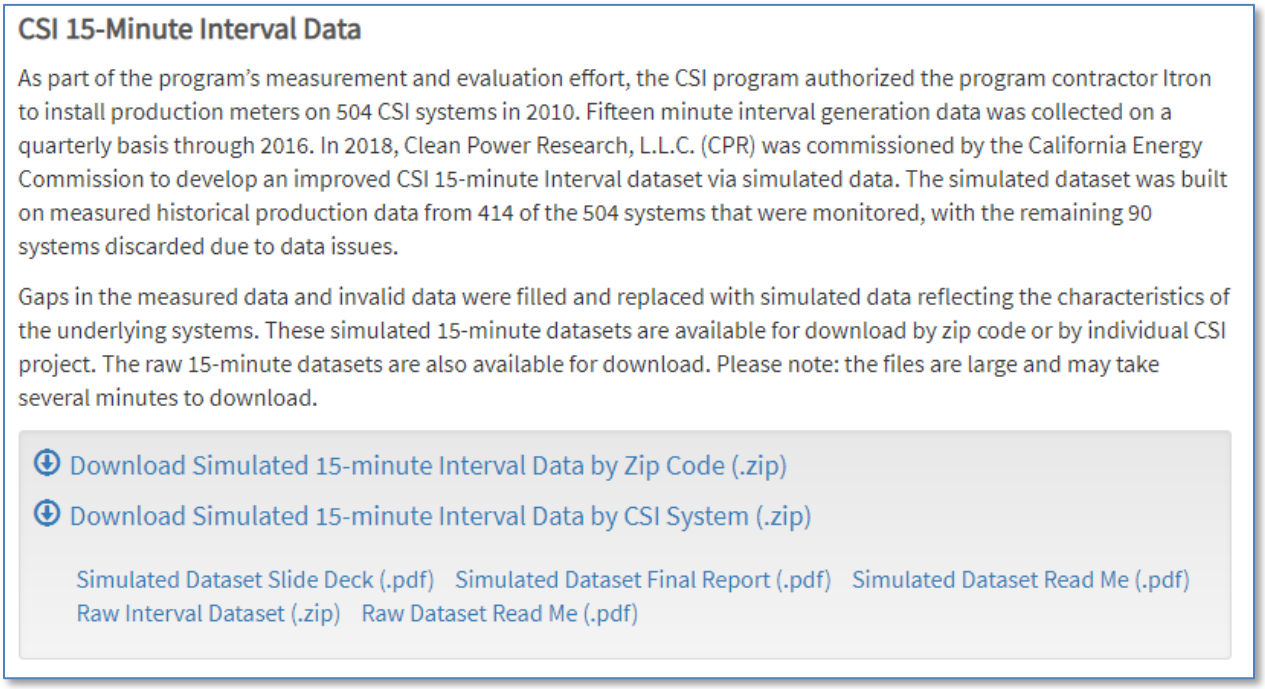

CaliforniaDGStats website hosts a key data deliverable from the project, including supporting information such as the relevant task report.

Source: CaliforniaDGStats (www.californiadgstats.ca.gov)

Additional Contacts

Itron sent the task report concerning California ISO integration to a number of ISO/RTOs and utilities. The report "Advanced Statistical Model Specifications for Incorporating On-Premise Solar PV Generation into a Short-Term Load Forecast Model" provides solar PV forecast improvement results and methods for assessing them. Recipients of the report include:

- Australian Energy Market Operator, Jack Fox
- Alberta Energy System Operator, Canada, Lars Renborg
- Bonneville Power Association, Justin Rohrboug
- California ISO, Amber Motley

- El Centro Nacional de Control de Energía, Mexico, Jesus Valenzuela Balbastro
- Electric Reliability Council of Texas, Calvin Opheim
- Independent Energy System Operator, Ontario Canada, Andrew Trachsell
- Independent System Operator New England, John Black
- Midcontinent Independent System Operator, Blogoy Borissov
- New York Independent System Operator, Arthur Maniaci
- Southwest Power Pool, Gunnar Shaffer
- Tennessee Valley Authority, Jeff House
- Western Power (now part of the Australian Energy Market Operator, David Weinholz

Ongoing Outreach

Ongoing outreach and future plans include:

- CPR is using improved forecasting methods derived from the project (Task 2, UCSD/SUNY Forecast Improvement) in its fleet forecasting for use by the California ISO, the IOUs, and out-of-state agencies. As CPR licenses its forecast software to new agencies (future licenses), embedded improvements will be available since they represent best practices for fleet forecasting.
- Itron will be able to better integrate the BTM fleet forecast into its neural network model using the methods developed by this project (Task 5, Short-Term Load Forecast) for the California ISO. This is an ongoing process between Itron and the California ISO that will directly benefit the California grid, and knowledge of forecast improvements will continue to spur technical advances.
- Data deliverables from Task 4, DER Production Database, will be usable by researchers, market analysts, policymakers, and others as a publicly available input data set. This data will continue to be offered for download for analysis where detailed 15-minute data is required.
- CPR will incorporate the spec inference engine into its SolarAnywhere® software to support PV asset managers. One company is already using this as a new service, which was not previously available to the industry. The application improves modeling by quantifying the impacts of shading and other real-world solar PV generation factors not typically available from system installers.

CHAPTER 10:
Benefits to Ratepayers

Lower Costs

Ratepayers benefit from an improved solar PV fleet production forecast because it will enable the California ISO to reduce net-load forecast uncertainty. This is because the cost of procuring both spinning and non-spinning reserves will be reduced, since fewer of these resources will be required.

If uncertainty is high, grid operators must readily make available generating units for quick action when required to meet demand. This calls for additional capacity, which is costly. In most cases, it also calls for fossil-fueled spinning reserves that burn a small amount of fuel even when units are off line. These fossil-fuel spinning resources can ramp up to full power rapidly if needed, reducing the time required to meet electricity load.

CPR developed forecast error valuation methodologies that will permit users to estimate the cost of improvements in forecast accuracy. These methods are also applied to forecast improvements developed in research from this project. Using the Resource Adequacy method, CPR estimates that the forecasting methods developed in this project will save California $176 million over the next 30 years. Over time, these forecast error valuation methodologies will assist in the long-term resource planning process and could ultimately reduce the need for planned resources.

Greater Reliability

The project improves solar PV production forecasts, with great potential to increase power system reliability across California.

Uncertainty in generation from both BTM and centralized solar PV generation imparts costs to ratepayers, because grid operators must carry spinning and non-spinning reserves to meet demand created by this uncertainty. Reducing uncertainty could potentially reduce the need for those required reserves.

With reduced uncertainty leading to a higher probability of meeting net loads, reliability will increase.

This project quantified specific reductions in reserve requirements. Project results will provide knowledge to grid operators and other California decision makers that will help the state plan for greater integration of DER generation while maintaining reliability standards. Based on the Resource Adequacy method, designed to eliminate 90 percent of forecast uncertainty, expected savings in the three-hour forecast would be $28 per kW of installed PV.

Environmental Benefits

Adoption of this project's results will improve California's environmental footprint, because greater DER generation will decrease use of fossil-fueled reserves, which would otherwise be required to accommodate PV forecast inaccuracy.

GLOSSARY AND LIST OF ACRONYMS

Term	Definition
Azimuth	The horizontal angular distance between the vertical plane containing a point in the sky and true north.
BTM	Behind-the-Meter: Generation connected on the customer side of the meter that impacts net load
California ISO	California Independent System Operator – the organization that manages the three IOU's electricity grid in California
Clear sky GHI	The theoretical GHI under clear sky conditions
Clear sky power	The theoretical power output of a PV system or fleet under clear sky conditions
CMV	Cloud Motion Vector: A forecast method in which the location of clouds in the future are calculated based on their current location and the prevailing wind speed and direction.
DAM	Day-Ahead Market
EPIC	Electric Program Investment Charge
FS	Forecast Skill. This is a measure of forecast accuracy.
GHI	Global Horizontal Irradiance: A measure of the sun's intensity as measured on a horizontal surface, including both the direct (beam) irradiance and the diffuse irradiance components.
GOES	Geostationary Operational Environmental Satellite: A satellite system, operated by the US National Oceanic and Atmospheric Administration, designed to support weather forecasting, severe storm tracking, and meteorology research
HASP	Hour-Ahead Schedule Process
Kt	Clearness index: percentage of clear sky, calculated by dividing GHI by clear sky GHI
LMP	Locational Marginal Price. This is the market price of power in a given hour at a given location.
MAE	Mean Absolute Error: A measure of forecast error based on the absolute difference in forecast versus measured value.
MAPE	Mean Absolute Percentage Error: A measure of forecast error based on the percentage difference in forecast versus measured value.
MBE	Mean Bias Error. This is a measure of forecast accuracy.

Term	Definition
Measured Load	The load seen at the customer meter, or the actual load minus any generation. For this report, this refers to the aggregate of all customer metered loads at the California ISO load zone
NNET	Neural Network. This is a machine learning approach used in load forecasting.
NWP	Numerical Weather Prediction: Any of several mathematical models of the atmosphere and oceans to predict the weather based on current weather conditions.
PG&E	Pacific Gas and Electric; the IOU that provides natural gas and electricity to much of Northern California
Reconstituted Load	Measured Load with behind-the-meter solar PV generation added to form an estimate of electricity consumption
rMAE	Relative Mean Absolute Error: A measure of forecast error based on the relative difference in the mean difference of forecast versus measured value.
RMSE	Root Mean Square Error
SCE	Southern California Edison; the IOU that provides electricity to much of Southern California outside of San Diego
SDG&E	San Diego Gas and Electric; the IOU that provides natural gas and electricity to San Diego and the surrounding area
SolarAnywhere	The Clean Power Research software family related to historical and forecast solar irradiance and solar PV energy system performance.
Solar PV	Solar Photovoltaic; a technology that uses semiconductors to convert solar irradiance into DC electrical power. This DC electrical power is usually converted to AC electrical power uses inverter(s).
Sc	Stratocumulus: A genus-type of clouds characterized by large dark, rounded masses, usually in groups, lines, or waves, the individual elements being larger than those in altocumulus, and the whole being at a lower height.
SURFRAD	The Surface Radiation budget (SURFRAD) network, developed for the US in the mid-1990s to support satellite system validation; numerical model verification; and modern climate, weather, and hydrology research applications.

REFERENCES

Akyurek, Bengu Ozge, Jan Kleissl. 2017. Closed-Form Analytic Solution of Cloud Dissipation for a Mixed-Layer Model. https://doi.org/10.1175/JAS-D-16-0303.1. Journal of Atmospheric Sciences 74, 2525–2556.

Arbizu-Barrena, Clara, José A. Ruiz-Arias, Francisco J. Rodríguez-Benítez, David Pozo-Vázquez, Joaquín Tovar-Pescador. 2017. Short-Term Solar Radiation Forecasting by Advecting and Diffusing MSG Cloud Index. https://doi.org/10.1016/j.solener.2017.07.045. Solar Energy 155, 1092–1103.

Beer, Charles G.P., Luna B. Leopold. 1947. Meteorological Factors Influencing Air Pollution in the Los Angeles Area. https://doi.org/10.1029/TR028i002p00173. Transactions American Geophysical Union 28, 173.

California DG Stats. 2018. https://www.californiadgstats.ca.gov/downloads/#_csi_15_id (accessed Jan. 1, 2018). California Public Utilities Commission.

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. California Energy Commission.

Chauvin, Rémi, Julien Nou, Stéphane Thil, Stéphane Grieu. 2016. Cloud Motion Estimation Using a Sky Imager. https://doi.org/10.1063/1.4949235. AIP Conference Proceedings.

Chow, Chi Wai, Bryan Urquhart, Matthew Lave, Anthony Dominguez, Jan Kleissl, Janet Shields, Byron Washom. 2011. Intra-hour Forecasting with a Total Sky Imager at the UC San Diego Solar Energy Testbed. https://doi.org/10.1016/j.solener.2011.08.025. Solar Energy 85, 2881–2893.

CIMIS. 2017. http://wwwcimis.water.ca.gov/Stations.aspx (accessed Dec. 7, 2017). California Irrigation Management Information System.

Clemesha, Rachel E.S., Alexander Gershunov, Sam F. Iacobellis, A. Park Williams, Daniel R. Cayan. 2016. The Northward March of Summer Low Cloudiness Along the California Coast. https://doi.org/10.1002/2015GL067081. Geophysical Research Letters 43, 1287–1295.

Coimbra, Carlos F.M., Jan Kleissl, Ricardo Marquez. 2013. Overview of Solar-Forecasting Methods and a Metric for Accuracy Evaluation. https://doi.org/10.1016/B978-0-12-397177-7.00008-5. Solar Energy Forecasting and Resource Assessment. pp. 171–194.

Crosbie, Ewan, Zhen Wang, Armin Sorooshian, Patrick Y. Chuang, Jill S. Craven, Matthew M. Coggon, Michael A Brunke, Xubin Zeng, Haflidi Jonsson, R. K. Woods, Richard C. Flagan, John H. Seinfeld. 2016. Stratocumulus Cloud Clearings and Notable Thermodynamic and Aerosol Contrasts Across the Clear–Cloudy Interface. https://doi.org/10.1175/JAS-D-15-0137.1. Journal of Atmospheric Sciences 73, 1083–1099.

Denholm, Paul, Kara Clark, Matt O'Connell. 2016. On the Path to SunShot. Emerging Issues and Challenges in Integrating High Levels of Solar into the Electrical Generation and Transmission System. https://doi.org/10.2172/1253978. National Renewable Energy Laboratory.

Dev, Soumyabrata, Florian M. Savoy, Yee Hui Lee, Stefan Winkler. 2017. Short-term Prediction of Localized Cloud Motion Using Ground-based Sky Imagers. https://doi.org/10.1109/TENCON.2016.7848499. IEEE Region 10 Annual International Conference, Proceedings/TENCON. pp. 2563–2566.

EPRI. 2017. Solar Power Forecasting for Grid Operations: Evaluation of Commercial Providers. EPRI Technology Insights, Nov. 2017.

GEBCO. 2017. General Bathymetric Chart of the Oceans. http://www.gebco.net/data_and_products/gridded_bathymetry_data/ (accessed Oct. 18, 2016).

Ghonima, M.S., Thijs Heus, Joel R. Norris, Jan Kleissl. 2016. Factors Controlling Stratocumulus Cloud Lifetime over Coastal Land. https://doi.org/10.1175/JAS-D-15-0228.1. Journal of the Atmospheric Sciences 73, 2961–2983.

Hartmann, Dennis L., Maureen E. Ockert-Bell, Marc L. Michelsen. 1992. The Effect of Cloud Type on Earth's Energy Balance: Global Analysis. https://doi.org/10.1175/1520-0442(1992)005<1281:TEOCTO>2.0.CO;2. Journal of Climate 5, 1281–1304.

Hahn, Carole, Stephen Warren. 2007. A Gridded Climatology of Clouds Over Land (1971–96) and Ocean (1954–97) From Surface Observations Worldwide. https://doi.org/10.3334/CDIAC/cli.ndp026e. Oak Ridge National Laboratory, Carbon Dioxide Information Analysis Center.

Huang, Hao, Jin Xu, Zhenzhou Peng, Shinjae Yoo, Dantong Yu, Dong Huang, Hong Qin. 2013. Cloud Motion Estimation for Short Term Solar Irradiation Prediction. https://doi.org/10.1109/SmartGridComm.2013.6688040. 2013 IEEE Int. Conf. Smart Grid Commun. SmartGridComm 2013 696–701.

Iacobellis, Sam F., Daniel R. Cayan. 2013. The Variability of California Summertime Marine Stratus: Impacts on Surface Air Temperatures. https://doi.org/10.1002/jgrd.50652. Journal of Geophysical Research: Atmospheres 118, 9105–9122.

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

ITRON. 2014. ITRON White Paper, Forecast Practitioner's Handbook: Incorporating the Impact of Embedded Solar Generation into a Short-Term Load Forecasting Model, (101376WP-01) 2014.

Jamaly, M., Jan Kleissl. 2012. Validation of SolarAnywhere Enhanced Resolution Irradiation in California. Intern. Rep.

Jimenez, Pedro A., Joshua P. Hacker, Jimy Dudhia, Sue Ellen Haupt, José A. Ruiz-Arias, Chris A. Gueymard, Gregory Thompson, Trude Eidhammer, Aijun Deng. 2016. WRF-SOLAR: Description and Clear-Sky Assessment of an Augmented NWP Model for Solar Power Prediction. https://doi.org/10.1175/BAMS-D-14-00279.1. Bulletin of the American Meteorological Society 97, 1249–1264.

Johnstone, James A., Todd E. Dawson. 2010. Climatic Context and Ecological Implications of Summer Fog Decline in the Coast Redwood Region. https://doi.org/10.1073/pnas.0915062107. Proceedings of the National Academy of Sciences of the U.S.A. 107, 4533–8.

Kankiewicz, Adam, john Dise, Elynn Wu, Richard Perez. 2014. Solar 2014: Reducing Solar Project Uncertainty with an Optimized Resource Assessment Tuning Methodology. 2014 American Solar Energy Society Annual Conf. 1–6.

Klein, Stephen A., Dennis L. Hartmann. 1993. The Seasonal Cycle of Low Stratiform Clouds. https://doi.org/10.1175/1520-0442(1993)006<1587:TSCOLS>2.0.CO;2. Journal of Climate.

Köhler, Carmen, Andrea Steiner, Yves-Marie Saint-Drenan, Dominique Ernst, Anja Bergmann-Dick, Mathias Zirkelbach, Zied Ben Bouallègue, Isabel Metzinger, Bodo Ritter. 2017. Critical Weather Situations for Renewable Energies – Part B: Low Stratus Risk for Solar Power. https://doi.org/10.1016/j.renene.2016.09.002. Renewable Energy 101, 794–803.

Lara-Fanego, V., José A. Ruiz-Arias, David Pozo-Vázquez, Francisco J. Santos-Alamillos, Joaquín Tovar-Pescador. 2012. Evaluation of the WRF Model Solar Irradiance Forecasts in Andalusia (southern Spain). https://doi.org/10.1016/j.solener.2011.02.014. Solar Energy 86, 2200–2217.

Lee, Jared A., Sue Ellen Haupt, Pedro A. Jiménez, M.A. Rogers, S.D. Miller, T.C. McCandless. 2017. Solar Irradiance Nowcasting Case Studies Near Sacramento. https://doi.org/10.1175/JAMC-D-16-0183.1. Journal of Applied Meteorology and Climatology 56, 85–108.

Mathiesen, Patrick, Craig Collier, Jan Kleissl. 2013. A High-resolution, Cloud-assimilating Numerical Weather Prediction Model for Solar Irradiance Forecasting. https://doi.org/10.1016/j.solener.2013.02.018. Solar Energy 92, 47–61.

Mathiesen, Patrick, Jan Kleissl. 2011. Evaluation of Numerical Weather Prediction for Intra-day Solar Forecasting in the Continental United States. https://doi.org/10.1016/j.solener.2011.02.013. Solar Energy 85, 967–977.

Miller, Steven D., Matthew A. Rogers, John M. Haynes, Manajit Sengupta, Andrew K. Heidinger. 2017. Short-term Solar Irradiance Forecasting via Satellite/Model Coupling. https://doi.org/10.1016/j.solener.2017.11.049. Solar Energy 168, 102-117.

Monforte, Frank, Christine Fordham, Jennifer Blanco, Stephan Barsun, Adam Kankiewicz, Ben Norris. 2016. Improving Short-term Load Forecasts by Incorporating Solar PV Generation. California Energy Commission, Contract Number: EPC-14-001 September 2016.

Monforte, Frank, J. Stuart McMenamin. 1998. Short-Term Energy Forecasting with Neural Networks. The Energy Journal, Volume 19, Number 4, 1998. International Association for Energy Economics.

NESDIS. 2017. GOES Imager Calibration. https://www.star.nesdis.noaa.gov/smcd/spb/fwu/homepage/GOES_Imager.php (accessed Jul. 17, 2017). National Environmental Satellite, Data and Information Service.

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

Peng, Zhenzhou, Dantong Yu, Dong Huang, John Heiser, Shinjae Yoo, Paul Kalb. 2015. 3D Cloud Detection and Tracking System for Solar Forecast Using Multiple Sky Imagers. https://doi.org/10.1016/j.solener.2015.05.037. Solar Energy 118, 496–519.

Perez, Richard, Pierre Ineichen, Kathy Moore, Marek Kmiecik, Cyril Chain, Ray George, Frank Vignola. 2002. A New Operational Model for Satellite-derived Irradiances: Description and Validation. https://doi.org/10.1016/S0038-092X(02)00122-6. Solar Energy 73, 307–317.

Perez, Richard, Adam Kankiewicz, James Schlemmer, Karl Hemker, Jr., Sergey Kivalov. 2014. A New Operational Solar Resource Forecast Service for PV Fleet Simulation. 40th IEEE PV Specialists Conference, 2014

Perez, Richard, Sergey Kivalov, James Schlemmer, Karl Hemker, David Renné, Thomas E. Hoff. 2010. Validation of Short and Medium Term Operational Solar Radiation Forecasts in the US. https://doi.org/10.1016/j.solener.2010.08.014. Solar Energy 84, 2161–2172.

Perez, Richard, James Schlemmer, Karl Hemker, Jr., Sergey Kivalov, Adam Kankiewicz, John Dise. 2016. Solar Energy Forecast Validation for Extended Areas & Economic Impact of Forecast Accuracy. 2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC).

Perez, Richard, James Schlemmer, Adam Kankiewicz, John Dise, Alemu Tadese, Thomas Hoff. 2017. Detecting Calibration Drift at Ground Truth Stations. A Demonstration of Satellite Irradiance Models' Accuracy. 2017 IEEE 44th Photovoltaic Specialist Conference (PVSC).

Rastogi, Bharat, Alton Park Williams, Douglas T. Fischer, Sam F. Iacobellis, Kathryn McEachern, Leila Carvalho, Charles Jones, Sara A. Baguskas, Christopher J. Still. 2016. Spatial and Temporal Patterns of Cloud Cover and Fog Inundation in Coastal California: Ecological Implications. https://doi.org/10.1175/EI-D-15-0033.1. Earth Interactions 20, 1–19.

Rienecker, Michele M., Max J. Suarez, Ronald Gelaro, Ricardo Todling, Julio Bacmeister, Emily Liu, Michael G. Bosilovich, Siegfried D. Schubert, Lawrence Takacs, Gi-Kong Kim, Stephen Bloom, Junye Chen, Douglas Collins, Austin Conaty, Arlindo da Silva, Wei Gu, Joanna Joiner, Randal D. Koster, Robert Lucchesi, Andrea Molod, Tommy Owens, Steven Pawson, Philip Pegion, Christopher R. Redder, Rolf Reichle, Franklin R. Robertson, Albert G. Ruddick, Meta Sienkiewicz, Jack Woollen. 2011. MERRA: NASA's Modern-Era Retrospective Analysis for Research and Applications. https://doi.org/10.1175/JCLI-D-11-00015.1. Journal of Climate 24, 3624–3648.

SolarAnywhere. 2017. Web-Based Service that Provides Hourly, Satellite-Derived Solar Irradiance Data Forecasted 7 Days Ahead and Archival Data Back to January 1, 1998. https://www.solaranywhere.com/products/solaranywhere-data/ (accessed Jul. 10, 2017).

Torregrosa, Alicia, Cindy Combs, Jeff Peters. 2016. GOES-derived Fog and Low Cloud Indices for Coastal North and Central California Ecological Analyses. https://doi.org/10.1002/2015EA000119. Earth and Space Science 3, 46–67.

Weinreb, Michael, Dejiang Han. 2011. Conversion of GVAR Infrared Data to Scene Radiance or Temperature. http://www.ospo.noaa.gov/Operations/GOES/calibration/gvar-conversion.html (accessed Jul. 10, 2017). NOAA Office of Satellite and Product Operations.

Williams, A. Park, Rachel Schwartz, Sam Iacobellis, Richard Seager, Benjamin I. Cook, Christopher J. Still, Gregory Husak, Joel Michaelsen. 2015. Urbanization Causes Increased Cloud Base Height and Decreased Fog In Coastal Southern California. https://doi.org/10.1002/2015GL063266. Geophysical Research Letters 42, 1527–1536.

Wood, Robert. 2012. Stratocumulus Clouds. https://doi.org/10.1175/MWR-D-11-00121.1. Monthly Weather Review 140, 2373–2423.

Yang, Handa, Jan Kleissl. 2016. Preprocessing WRF Initial Conditions for Coastal Stratocumulus Forecasting. https://doi.org/10.1016/j.solener.2016.04.003. Solar Energy 133, 180–193.

Yang, Handa, Ben Kurtz, Dung Nguyen, Bryan Urquhart, Chi Wai Chow, Mohamed Ghonima, Jan Kleissl. 2014. Solar Irradiance Forecasting Using a Ground-based Sky Imager Developed at UC San Diego. https://doi.org/10.1016/j.solener.2014.02.044. Solar Energy 103, 502–524.

Zhong, Xiaohui, Dipak K. Sahu, Jan Kleissl. 2017. WRF Inversion Base Height Ensembles for Simulating Marine Boundary Layer Stratocumulus. https://doi.org/10.1016/j.solener.2017.02.021. Solar Energy 146, 50–64.