Development of Flexible Ramp Product Procurement for the California ISO using Probabilistic Solar Power Forecasts

Benjamin F. Hobbs[1], Venkat Krishnan[2], Jie Zhang[3], Hendrik F. Hamann[4], Carlo Siebenschuh[4], Rui Zhang[4], Binghui Li[3], Ibrahim Krad[2], Evangelia Spyrou[2], Yijiao Wang[1], Shu Zhang[1]

[1]Robert O'Connor Sustainable Energy Institute, Johns Hopkins University, Baltimore, MD (USA)

[2] Energy Systems Integration Directorate, National Renewable Energy Laboratory, Golden, CO (USA)

[3] Department of Mechanical Engineering, The University of Texas at Dallas, Richardson, TX (USA)

[4] IBM Thomas J. Watson Research Center, Ossining, NY (USA)

Abstract

How can probabilistic solar forecasts lower costs and improve reliability for independent system operator (ISO) markets? We tackle this question in three steps. First, we enhance an existing solar forecasting system to provide well-calibrated hours-ahead probabilistic forecasts. We then relate the degree of uncertainty in those forecasts to error distributions for net load ramps for the California ISO (CAISO) using statistical and machine learning methods. Projected net load errors conditioned on solar uncertainty are translated into flexible ramp requirements that therefore reflect real-time meteorological and solar conditions, improving on typical ISO procedures. Finally, a multi-period look-ahead production cost model quantifies how conditional ramp requirements can a) decrease operating costs by lowering requirements compared to often conservative unconditional methods, and b) reduce generation scarcity events and consequently improve reliability by increasing flexibility requirements at times when unconditional forecast-based requirements understate actual ramp uncertainty.

Key Words: Probabilistic solar power forecasts, flexible ramp product requirements, California, production cost, reliability, machine learning.

1. Introduction and Background

As renewable penetration grows in California and elsewhere, both in the form of behind-the-meter rooftop solar, and in the form of grid-scale wind and solar facilities, the short-term variability in system net load that must be met by dispatchable thermal, hydro, and increasingly battery resources is growing (Mills et al., 2021). The CAISO procures reserves of various types to accommodate this variability for the California market on various time scales, ranging from 30-minute non-spinning reserves, 10-minute spinning reserves, regulation-up, and regulation-down, with the last two handling unexpected variations on a sub-5 minute scale. These are acquired in the day-ahead CAISO market, with adjustments possible in the real-time market (part of the west-wide Energy Imbalance Market (EIM), whose other entities only procure or sell real-time energy in the EIM). In addition, the CAISO has developed a new real-time product in 2016 called the flexible ramp product (FRP). It is designed to position resources in the 15- and 5-minute intervals of the EIM to feasibly accommodate unexpected deviations in net load ramp from interval to interval either in the upward or downward direction, and to compensate resources for any resulting foregone energy revenues (Wang and Hobbs, 2015; CAISO, 2020). Finally, the CAISO is proposing a day-ahead product called the imbalance reserve product that will not only procure resources that could meet the real-time FRP needs, but also accommodate deviations between day-ahead and real-time forecasts in overall net load with a predetermined reliability of 95% (Angelidis, 2020).

In deciding how much flexible ramp product to procure, the CAISO considers the historic (approximately previous month) of distribution of ramp forecast errors in the relevant time interval. By describing a histogram of those errors, the 2.5% and 97.5% percentiles can be estimated, and used to define the MW of "uncertainty component" in the down- and up-directions to be acquired. The net down- and up-requirement for ramp in a given period is then defined as the expected ramp forecast plus those two uncertainty components, respectively (which are usually negative and positive, in turn) (Fig. 1).

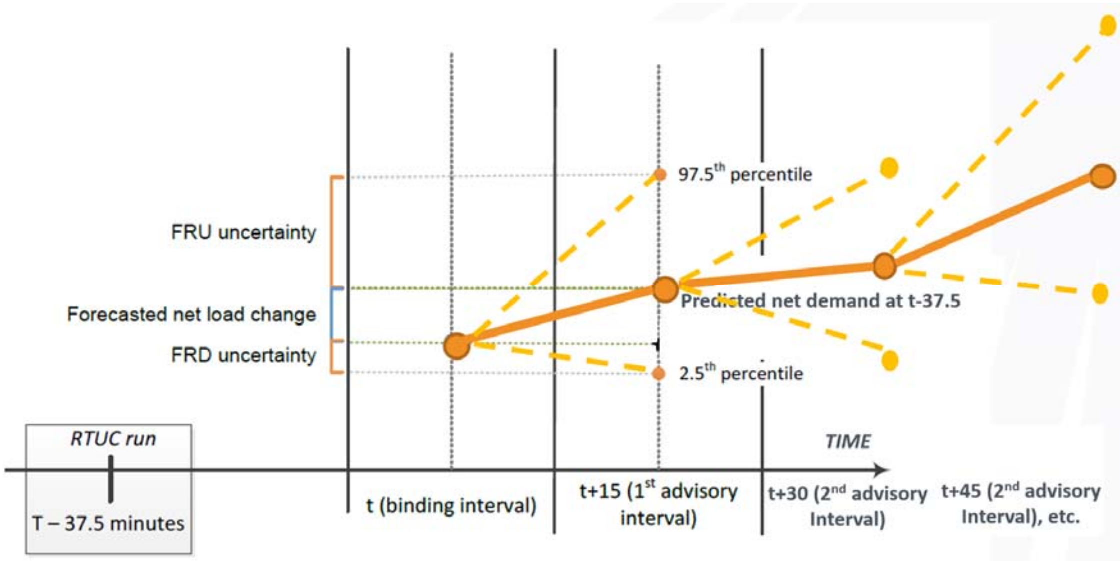

Fig. 1: Schematic of definition of uncertainty and forecast net load change components of CAISO flexible ramp product requirements (FRU = flexible ramp up, FRD = flexible ramp down), based on 2.5th and 97.5th percentiles of historical ramp forecast error. Note that the requirements area defined for all intervals in the 15-minute real-time market (binding interval, and subsequent advisory intervals in the multi-interval optimization). Source: Adapted by authors from CAISO.

These requirements are then translated into demand curves for MW of FRP in the up and down product in each interval that reflect the understanding that the incremental value of FRP declines as more is procured, but that there is not a fixed requirement or threshold below which a product has a high value (reflected in a large violation penalty in the market software) and above which it has no value. Further, the CAISO recognizes, as does most of the industry (Costilla-Enriquez et al., 2021), that defining requirements independent of information on weather and renewable energy production conditions will result in, some cases, overly conservative requirements well in excess of the actual need in a given day, thus inflating costs. While in other cases when uncertainty is greater, procurement will be too little, exposing the system to risks of inadequate flexibility reserves and undesirably high probabilities of load balance violations. Since solar generation variability is a major source of net load uncertainty, it is logical to expect that forecasts of the net load ramp uncertainty components shown in Fig. 1 could be usefully conditioned on weather and renewable conditions (CAISO, 2020), especially solar uncertainty.

In this paper, we summarize the procedures and results of research directed at translating probabilistic solar forecasts into weather-conditioned projections of FRP needs, and the subsequent production cost-based comparison of those revised requirements with the present CAISO unconditioned histogram approach for the western US markets (Fig. 2). The rest of the paper is organized into three parts: probabilistic solar forecast development (Section 2); 2) relating forecast uncertainty to CAISO real-time load ramps, yielding solar uncertainty-conditioned ramp requirements (Section 3); and 3) production cost-based assessment of cost savings and reliability improvements using a model system based upon the CAISO and western US power markets (Section 4).

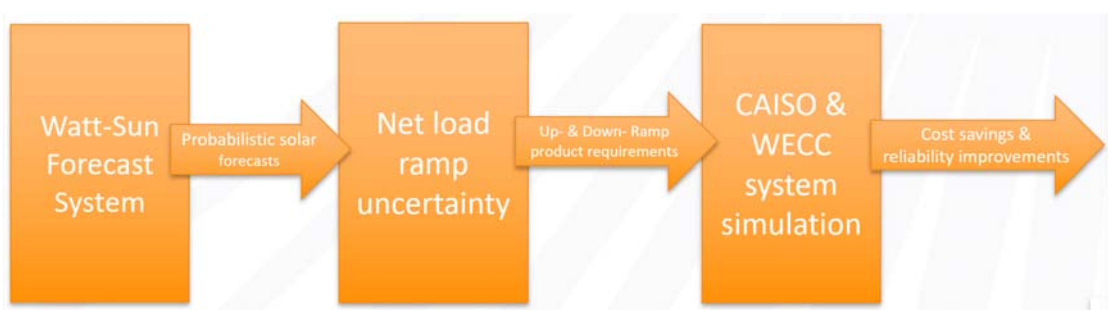

Fig. 2: Organization of analysis of solar forecast-based FRP requirements

2. Probabilistic solar forecasting

We have taken advantage of IBM's big-data platform Physical Analytics Integrated Repository and Services (PAIRS) to generalize the Watt-Sun solar forecasting system (Hamann et al. 2017) to generate probabilistic forecasts. PAIRS curates terabytes of numerical weather prediction (NWP) models and historical data, including the high-resolution GOES-R imagery (NOAA L2 product's cloud optical depth). By blending multiple NWPs and imagery data using deep learning methods and quantile regressions to obtain a set of critical percentiles, we developed a big data-driven probabilistic forecasting system, whose flow chart is shown in Fig. 3. The system has been implemented for 10 observation sites each in the CAISO and Mid-Continent ISO footprints for GHI forecasts. A prototype raster-based system using the GOES-R imagery has also been developed to create GHI irradiance forecasts on a 3 km pixel grid covering locations without direct observations.

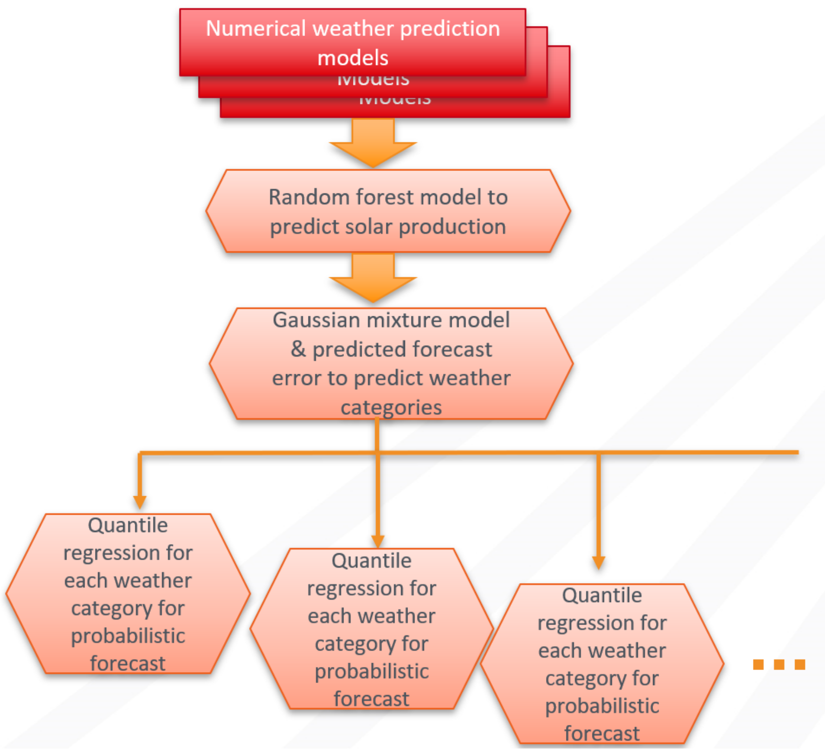

Fig. 3: Flow chart showing Probabilistic Watt-Sun development of probabilistic forecasts of solar irradiance (GHI)

Fig. 4 shows example Watt-Sun probabilistic forecasts, illustrating sunny and cloudy days. Based on P-P plot scores, the Watt-Sun probabilistic forecasts are better calibrated than baseline persistence and High-Resolution Rapid Refresh (HRRR) bias-corrected forecasts (Fig. 5). A P-P score measures the deviation (mean absolute value) of a plot of predicted versus empirical error cumulative distributions from the perfect (45^{o}) calibration line.

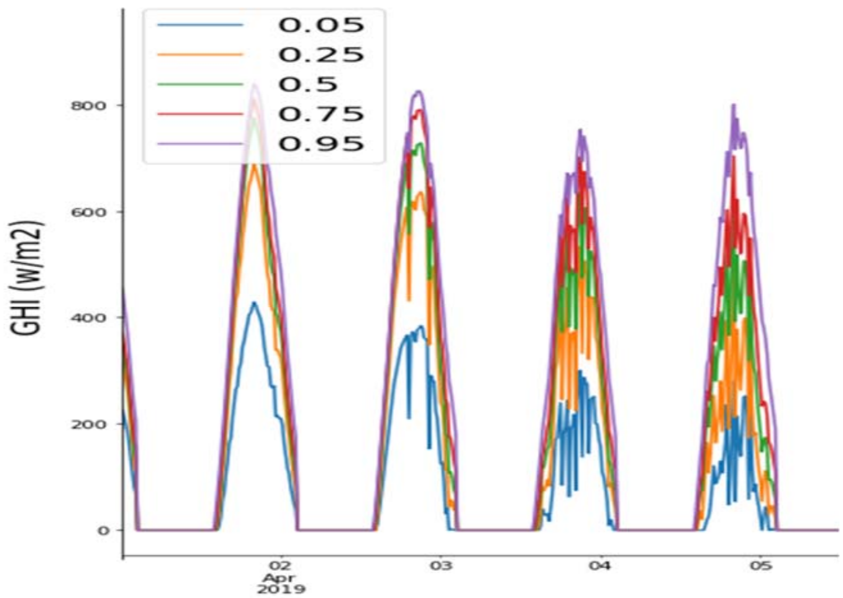

Fig. 4: Site-specific probabilistic global horizontal irradiance (GHI) forecast (April 2-5, 2019, Topaz site, 5th, 25th, 50th, 75th, and 95th percentiles), Probabilistic Watt-Sun 1.0 system

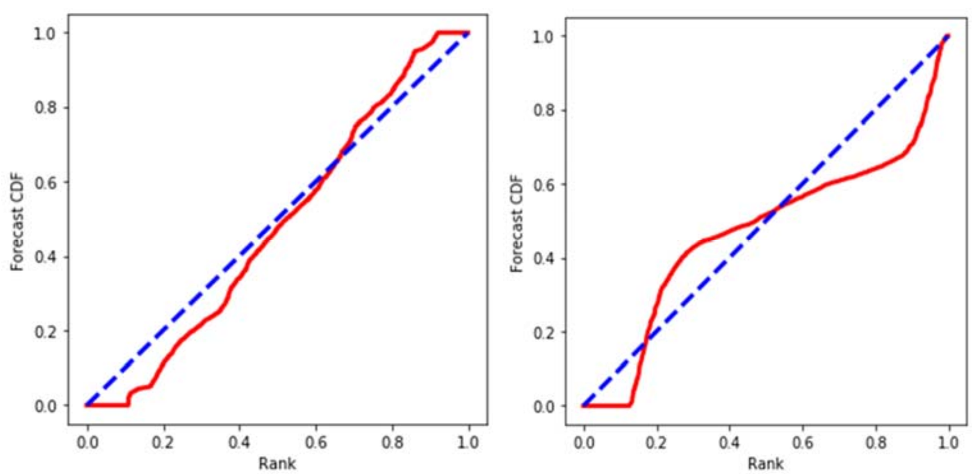

Fig. 5: Example P-P for Probabilistic Watt-Sun 1.0 forecast (P-P Plot score=0.054) and HRRR bias-corrected forecast (P-P Plot score=0.086)

3. Using probabilistic solar forecasts to create solar-conditioned ramp requirements

As summarized in Section 1, above, in 2016, the CAISO introduced ramp products in its real-time markets that procure generation "ramping" capacity so that potential real-time market ramps due to forecast uncertainties can be managed feasibly, in order to reduce the frequency of generation scarcity events and real-time price spikes. Specifically, the up- and down-flexible ramp product (FRP) addresses both expected 5-minute ramps plus an uncertainty component representing possible positive and negative errors, respectively, in net load forecasts, in particular the difference between so-called binding interval and subsequent first advisory interval forecasts (so as to procure enough ramping flexibility for immediate future market clearing intervals). Presently, to determine the uncertainty component, the CAISO uses histograms of net load forecast errors on an hourly basis, considering all errors in that hour from the previous 40 days (if the day is a weekday) or 20 days (if the day is a weekend) (www.caiso.com). After the construction of histograms, the upper and lower bounds of the 95% confidence interval are used as the up- and down-FRP requirements, respectively, for that hour. The 95% confidence interval for accounting for uncertainties is a well-accepted industry standard to strike a balance between reliability and economics. We term these "unconditional" or "solar independent" FRP requirements.

Ideally, the forecast error tends to be greater during a partially cloudy day while being lower in a sunny or completely overcast day. However, by constructing histograms of forecast errors purely from historical data, the CAISO's unconditional method does not reflect the latest weather information and often leads to overestimation of ramp uncertainty under sunny weather conditions and underestimation under cloudy conditions. As clearly depicted in Fig. 6, below, CAISO's requirements in two close days (8/7/19 and 8/12/19) are almost identical since the histograms are similar. However, as the purple lines in the figure show, the weather conditions in these two days are totally different: the left figure presents a lower and fluctuating profile of system-wide solar power, suggesting relatively cloudy weather, while the right profile shows a smooth and stable output, consistent with sunny conditions. Consequently, the actual FRP requirements should have been significantly different: greater uncertainties in cloudy days should contribute to higher FRP requirements, while sunny weather should be a reason for lower FRP requirements due to reduced uncertainty. If the same amount of FRP is procured in both days, the cloudy day would see a shortage of FRP, hence a compromised reliability level, while the market cost-effectiveness would be reduced in the sunny day due to over-procurement.

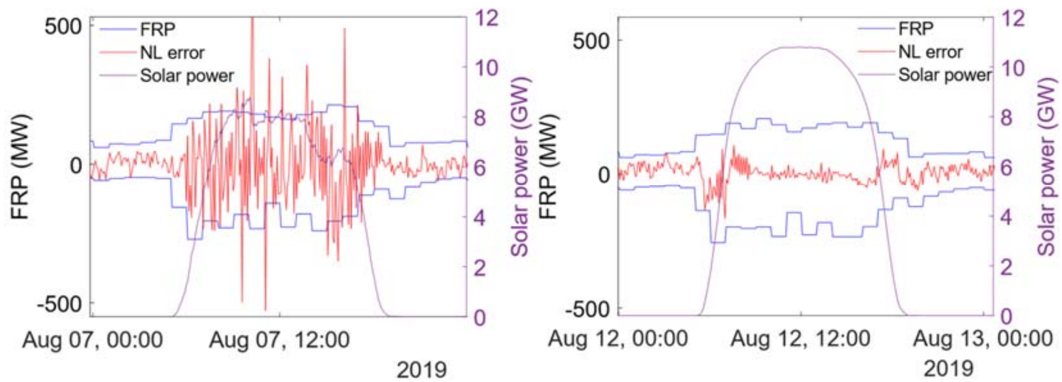

Fig. 6: Realized net load forecast errors (red) compared with CAISO's requirements (blue) from two days in August 2019 (cloudy (left) and sunny (right)). The purple curves show the CAISO's real-time forecasts of system-wide solar power production. (Time is UTC)

In order to improve current industry practice, we used quantile regression methods (e.g., Fig. 7) and machine learning methods to relate the 95% confidence intervals in ne load forecast errors to uncertainty in forecasted GHI in order to create solar-conditioned FRP requirements. Fig. 7 illustrates that there is a strong dependence of up-ramp uncertainty (as measured by the spread of values above zero) on the width of the 25th-75th percentile prediction interval from the Watt-Sun probabilistic forecasting system.

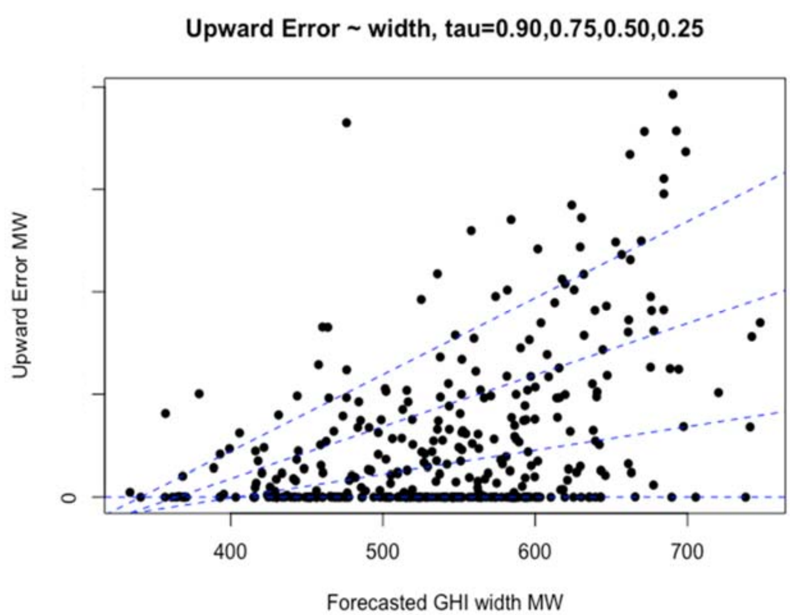

Fig. 7: 11 a.m.-2 p.m. May 2019: Quantile regression results for upward CAISO errors for (overestimated) real-time load forecasts as function of GHI 50% confidence interval width (from top, 90th, 75th, 50th, and 25th percentiles of load forecast errors)

Various specifications of the machine learning approach (based on the kth-nearest neighbor (kNN) classification method) were also tested with different combinations of GHI variables (e.g., median, 50% confidence interval width, volatility in 50% prediction intervals from 15-minute to 15-minute market interval) at various sites. Fig. 8 compares the out-of-sample performance of the kNN-based method for estimating ramp requirements relative to the unconditional method for February 2020, which are trained with data starting in January 2020. Displayed there are trade-offs between reliability levels and oversupplies in the form of Pareto frontiers for the case of 2-dimensional classifiers. The point at the intersection of two dashed lines in the figure represents the baseline using the CAISO's unconditional histogram method. The two dashed lines divide the plane into four quadrants (SW, SE, NE, and NW), where points in the southwest quadrant (SW) indicate improvements (reductions) in both shortage of reserves (an index of reliability) and oversupply, hence better solutions than the original implementation. The colored lines represent the performance of machine-learning based models derived from data from different solar production sites in the CAISO. A set of points for each site results from adjusting the parameters of each model to yield more (upper left) to less (lower right) conservative requirements, where "conservative" means that larger requirements and therefore lower probability of actual ramp errors lying outside the requirement, but higher procurement cost. In Fig. 8, a significant percentage of kNN points fall in quadrant SW, suggesting that the kNN-based method can both reduce average procurement requirements while improving average reliability levels. For instance, the figure shows that the best Pareto curve can reduce FRP oversupply to meet the present reliability performance (just under 8%) by a fifth (from 325 to 250), where "oversupply" is measured by the amount of FRP in excess of the actual amount needed.

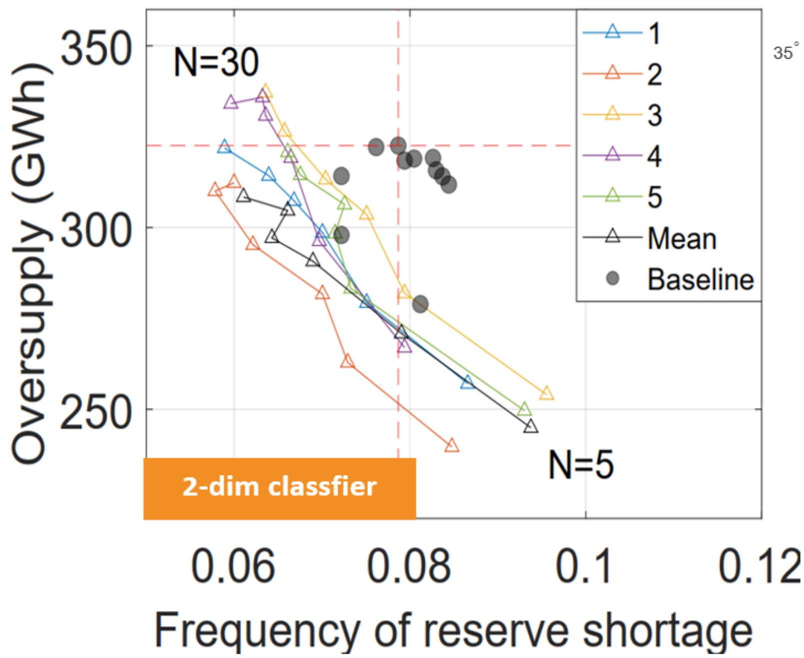

Fig. 8: Pareto diagram showing performance of kNN-based machine learning requirements for up-FRP based on probabilistic solar forecasting data from each of five solar production sites in the CAISO (labelled as sites 1-5, and also including the average performance). Performance is compared to the CAISO baseline histogram method, which is not conditioned on weather. The markets represent performance under alternative estimations from more conservative (upper left, showing more requirements and lower likelihood of reserve shortage, resulting from using larger training sets, starting from N=30 previous days at the extreme upper) to less conservative (lower right, resulting from using smaller training sets, starting from N=5 at the extreme lower point).

Ultimately the solar forecast-based requirements used in the simulations of the next section consisted of the kNN results based on the latest weather information. To select the optimal parameters for the kNN-based method–i.e., classifiers and the number of nearest neighbours–we designed a multi-objective optimization algorithm to dynamically select the parameters (Li et al., 2021).

Fig. 9 gives an example of a set of FRP-up and -down requirements for the CAISO resulting from this method (blue) that were used in the simulation of the next section, contrasting them to the FRP-up requirements based upon the CAISO baseline unconditional histogram method (orange). By definition, the requirements differed only during day-light hours. The results show that FRP-up requirements are increased during late afternoon hours to enhance system reliability, while FRP-down requirementts are reduced throughout the day. This suggests that the probabilistic solar forecasts are consistent with greater risks of decreases in solar production later in the day, but show smaller than typical risks of increased production throughout daylight hours. We next describe our analysis of the implications of these changed requirements for system production costs and FRP procurement costs.

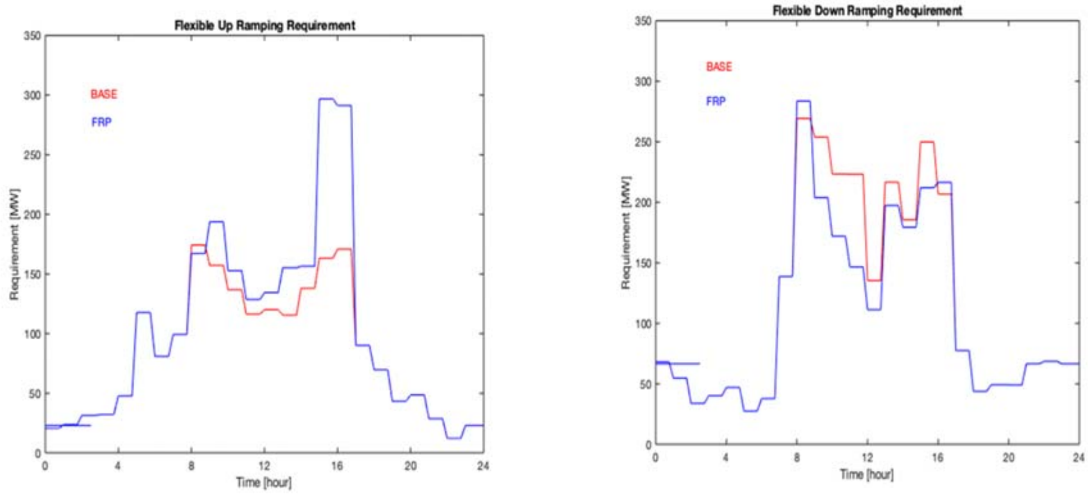

Fig. 9: FRP-up requirements defined by the CAISO baseline unconditional method and by the proposed probabilistic solar forecast conditional approach for example day during March 2020.

4. Production cost assessment of benefits of solar forecast-based flexible ramp product requirements

Our simulation of the benefits of solar-conditioned FRP requirements considered the IEEE 118 bus reliability test system with modifications reflecting California ISO conditions, modified to reflect high renewable penetration (nameplate solar penetration of approximately 70%, with a maximum instantaneous penetration exceeding 85%). Due to data and high-performance computing facility availability, we focused on operations during a few selected days (March 9-30, 2020).

Three scenarios were considered in this analysis. Scenario 1 is the 'baseline' simulation which captures the current unconditional FRP methodology, using available historical data from the CAISO. Scenario 2 is the 'conditional FRP' scenario in which FRP requirements are updated based on the methods outlined above to develop more efficient ramping requirements based on probabilistic solar power forecasts. Scenario 3 is the 'perfect' forecast scenario which models system operations if the operator has perfect knowledge of the system (i.e., zero netload forecast errors). The difference between Scenario 3 and any of the other scenarios' production cost gives a quantitative measure of "*uncertainty induced*" costs (which are a function of generation dispatch change-related costs to account for uncertainties and FRP procurements, as well as any costs for not meeting the FRP requirements as reflected in energy price spikes due to ramping scarcity).

All the analysis was performed using the FESTIV (Flexible Energy Scheduling Tool for Integrating Variable generation) (Ela et al., 2019), a multitemporal production cost model jointly developed by NREL and the Electric Power Research Institute. Table 1 shows that the conditional FRP, based on probabilistic solar forecasts, is highly successful, reducing uncertainty-related costs by almost 90% (from $443K to $48k) in a 22-day period in March 2020. Most of these benefits occurred in a one-week period mid-month (Fig. 10).

Tab. 1: Summary of FESTIV simulation results for CAISO-like IEEE 118 Bus, March 9-30, 2020

Index \ Scenario:	1 (Baseline)	2 (Conditional/Solar Informed FRP)	3 (Perfect Info)
Production Cost [$M]	23.445	23.049	23.002
Uncertainty Cost [$M]	0.443	0.048	- - -
Renewable Curtailment [GWh]	69.1	63.9	63.3

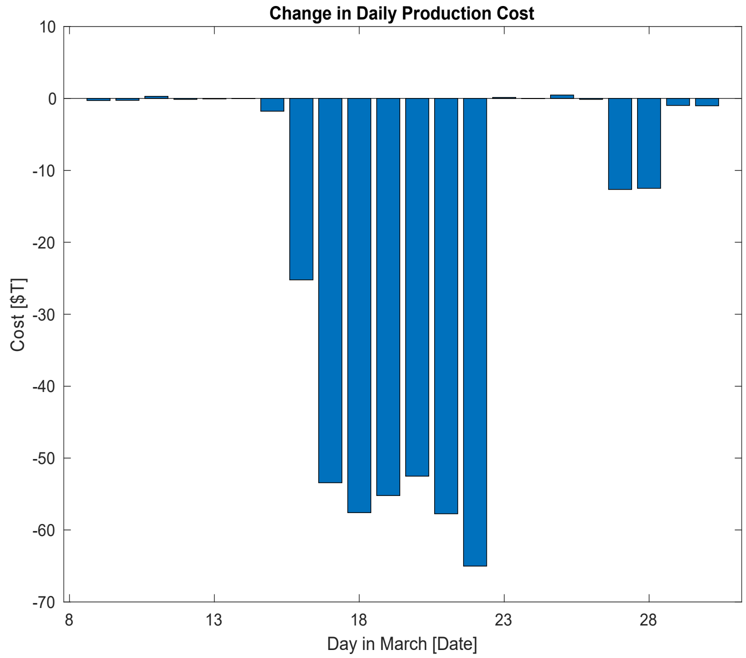

Fig. 10: Distribution of production cost reductions from using solar forecast-conditioned FRP requirements for 22 sample days in March 2020 compared to the CAISO baseline method, based upon 118-bus IEEE RTS 'CAISO-like" system

The benefits were due to reduction of FRP requirements from the solar-conditioned requirements method for some intervals, leading to reduced commitment of conventional generation and turning off of large generators with spinning capacity. This enabled the system to accommodate more solar output and reduce curtailment. Fig. 11 illustrates this impact by comparing the dispatch of four particular thermal generators in the evening of March 16, 2020, under the baseline and conditional FRP requirements scenarios (left and right sides of figure, respectively). The blue areas are the real-time energy production, and the orange areas are the "head room" that can accommodate unexpected upward ramps in the net load. The baseline scenario (left) has relatively low evening FRP-up requirements, so the day-ahead market believes that it has sufficient excess undispatched capacity on-line to meet the solar to thermal handoff in the evening. In real-time, however, the ramp-up is higher than expected, and this is reflected in the small amounts of "head room" (orange areas) in the second and third 15-minute intervals in the left side of Fig. 11. The real-time market detects that available capacity is depleted in those intervals, and so a larger (fifth) generator is then committed in the real-time market's short-term unit commitment procedure, and the four generators back down slightly in the fourth and fifth intervals to accommodate the minimum operating capability of the newly committed generator. By contrast, the larger FRP-up requirements in the solar-conditioned calculations (e.g., Fig. 10) force commitment in the day-ahead market of an additional smaller generator outside of the set of four generators considered in Fig. 11. This commitment leaves sufficient headroom to accommodate the larger than expected up-ramp in net load in real-time. The extra headroom shown as the orange areas in the second and third intervals in the right-side figure eliminate the need to commit the larger (fifth) generator we committed in the baseline scenario. The result is significant cost-savings, as shown for this day in Fig. 9.

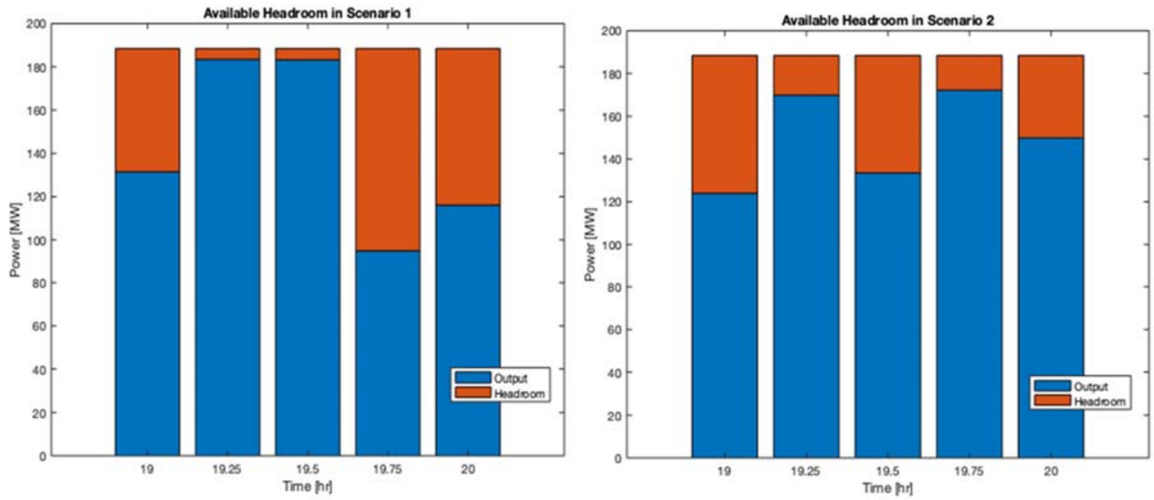

Fig. 11: Real-time dispatch of four impacted generators on March 16, 2020, under baseline FRP requirements (left) and solar conditioned requirements (right), illustrating that insufficient FRP-up in baseline case resulted in inadequate head-room in intervals two and three, necessitating real-time commitment of a costly fifth large generator in interval four.

Reliability improved in some other intervals under the solar-conditioned FRP requirements by procuring more FRP and avoiding generation scarcity, particularly during the evening ramping intervals as solar production tapers off. In other days, smaller generators were committed by the day-ahead model, which turned out to yield a lower minimum thermal production (in particular, a lower sum of Pmin levels for on-line generators) (e.g., Fig. 12). This allowed the system to deploy rather than curtail solar production in some hours, saving production costs.

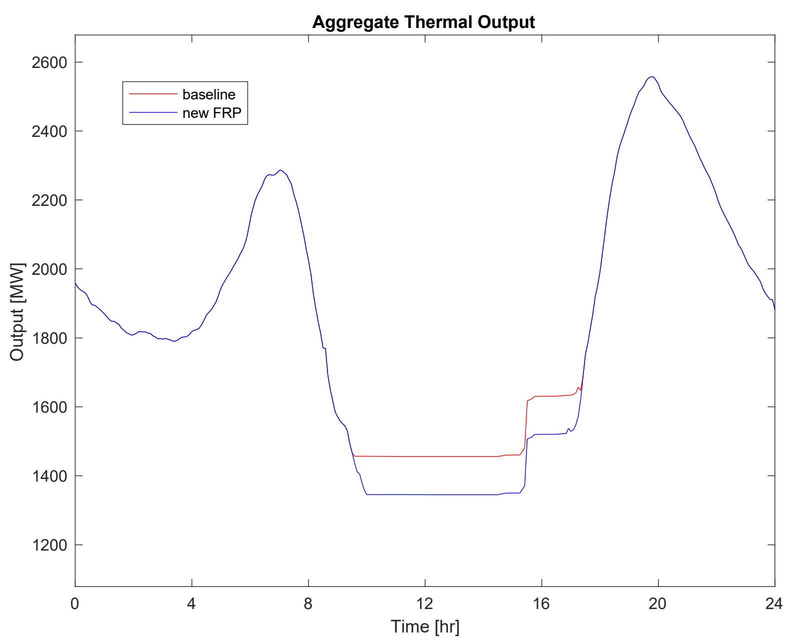

Fig. 12: Actual dispatch of all thermal generation in "CAISO-like" 118 bus IEEE TRS system under baseline (unconditional) FRP requirements (red) and solar conditional requirements (blue). Reduced thermal production the solar conditioned case accommodated increased utilization (reduced curtailment) of solar output.

We note that systems of different sizes or with more or less generation flexibility could change the above quantitative conclusions, although the qualitative improvements seen by the use of solar conditioned FRP requirements are likely to still apply.

5. Conclusion

As variable renewables increase their penetration, the need for—and cost of—operating reserves grows. Also, therefore, the potential benefits of conditioning reserve requirements on weather and renewable operating conditions will increase as well. In this paper, we have used probabilistic solar forecasts to definition weather-conditioned requirements for the California market's newest type of operating reserves, called the flexible ramp product. We summarize a k^{th}-nearest neighbor-based methodology relating solar forecast uncertainty (prediction interval width) to upward and downward uncertainty in 15-minute ramps in the California Independent System Operator's real-time market. Simulations of a power system whose generator mix characteristics are similar to California's illustrate the potential cost savings of using solar uncertainty-conditioned forecasts, which arise from more efficient unit commitment and reduced solar curtailment.

6. Acknowledgments

This material is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under Solar Energy Technologies Office (SETO) Agreement Number EE0008215. We thank the following Independent System Operator staff for their advice and data provided over the course of the project: Amber Motley, Guillermo Bautista-Alderete, Clyde Loutan, Rebecca Webb, and Hong Zhou of the California ISO, and Stephen Rose and Blagoy Borissov of the Mid-Continent ISO. Guidance from DOE SETO staff, especially John Seuss and Kemal Celik, are gratefully acknowledged. However, the authors are solely responsible for any opinions or errors contained in this paper.

7. References

Angelidis, G., 2021. Day-Ahead Market Enhancements, Appendix B: Draft Technical Description, Vers. 6.1. Report, California ISO, Aug. 6, www.caiso.com/InitiativeDocuments/AppendixB-DraftTechnicalDescription-Day-AheadMarketEnhancements.pdf (Last referenced Oct. 10, 2021)

California ISO, 2020. Flexible Ramping Product Refinements Initiative Appendix C - Quantile Regression Approach, Report, May 8, www.caiso.com/InitiativeDocuments/AppendixC-QuantileRegressionApproach-FlexibleRampingProductRequirements.pdf (Last referenced Oct. 10, 2021)

California ISO, 2021. Final Proposal: Flexible ramping product refinements, August 31, www.caiso.com/InitiativeDocuments/FinalProposal-FlexibleRampingProductRefinements.pdf (Last referenced Oct. 10, 2021)

Costilla-Enriquez, N., Ortega-Vazquez, M., Tuohy, A., Motley, A. and Webb, R., 2021. Operating Dynamic Reserve Dimensioning Using Probabilistic Forecasts, Report, Electric Power Research Institute, https://arxiv.org/pdf/2108.09362.pdf (Last referenced Oct. 10, 2021)

Ela, E., Palmintier, B. and Krad, I., 2019. FESTIV (Flexible Energy Scheduling Tool for Integrating Variable generation). National Renewable Energy Lab. (NREL), Report, Golden, CO, www.osti.gov/biblio/1504740 (Last referenced Oct. 10, 2021).

Hamann, H.F., et al., 2017. A multi-scale, multi-model, machine-learning solar forecasting technology, IBM, Final Technical Report, February, www.osti.gov/servlets/purl/1395344/ (Last referenced Oct. 10, 2021)

Li, B., Feng, C., Siebenschuh, C., Zhang, R., Spyrou, E., Krishnan, V., Hobbs, B.F., Zhang, J. 2021. Sizing ramping reserve using probabilistic solar forecasts: A data-driven method. Working Paper, University of Texas, Dallas.

Mills, A.D., Seel, J., Millstein, D., Kim, H., Bolinger, M., Gorman, W., Wang, Y., Jeong, S. and Wiser, R.H., 2021. Solar-to-Grid: Trends in System Impacts, Reliability, and Market Value in the United States with Data through 2019. Lawrence Berkeley National Laboratory, Berkeley, CA, https://emp.lbl.gov/sites/default/files/solar-to-grid_technical_report.pdf

Wang, B. and Hobbs, B.F., 2015. Real-time markets for flexiramp: A stochastic unit commitment-based analysis. IEEE Transactions on Power Systems, 31(2), pp.846-860.