

The value of Solar Power Forecasts – Applied Performance Metrics

Richard Perez¹, Jim Schlemmer¹, Sergey Kivalov¹, Karl Hemker, Jr.¹, Adam Kankiewicz², John Dise², and Alemu Tadese²

¹ Atmospheric Sciences Research Center, SUNY, Albany, New York, 12203, USA

² Clean Power Research, Napa, California, 94558, USA

Abstract: We introduce a new class of applied economic metrics to quantify the accuracy of solar power forecast in an operational context. These metrics are derived from the nominal cost of energy buffers that would be sufficient to make up for any forecast error so that, in effect, the output of a PV fleet operated in synergy with such energy buffers can be forecasted with 100% certainty. Two operational examples of this new class of metric are presented: (1) where forecasts are applied to predict solar production at an arbitrary time ahead with an arbitrary temporal or spatial granularity; and (2) where forecasts are applied to mitigate output fluctuations and ramps at any temporal or geographic scale. We also show that, in addition to evaluating the performance of forecast models, these metrics can also be applied directly to evaluate the cost of, and facilitate high PV penetration.

Index Terms —solar forecast, solar resource, metrics, high solar penetration, variability mitigation, forecast certainty

I. INTRODUCTION

The MAE and other standard accuracy metrics such as the Root Mean Square Error, the hit/miss ratio, or probabilistic forecast confidence intervals, are useful indicators of a model's performance. However these indicators are statistical measures that imply a non-zero probability that some forecasts will be missed. Here we turn the question around by determining the cost of the hardware sufficient to transform an imperfect forecast into a perfect forecast.

We examine two applications where forecasts would be operationally applied to facilitate the integration of PV fleets onto the power grid:

1. Power output guaranties for energy markets: The cost of missed forecasts depends on the rules and regulations in effect in a particular energy market, rules that are intended to account for the costs of increased reserves and grid flexibility necessary to manage PV output uncertainty. Fundamentally however, the cost of missed forecasts is traceable to two underpinnings: (1) procuring the missing energy in case of forecast overestimation, and (2) absorbing excess energy in case of forecast underestimation. This is illustrated in Figure 1at left. These tasks can be achieved using a buffer storage system. The buffer storage nominal price per unit of PV is the metric: the better the forecast, the smaller the buffer.
2. Ramp rate mitigation. A recent study by the authors showed that ramp rates can be mitigated to any desired maximum level at any time scale by applying an optimally sized energy buffer. The study also showed that for any given ramp mitigation objective, the buffer size could be considerably reduced if power output is known in advance, i.e., if it is perfectly forecasted. This is shown in Figure 1 at right. The metric in this case is defined by the storage buffer size needed with a given forecast in relation to the cases of no forecast (i.e., persistence) and ideal forecast.

II. METHODOLOGY: METRIC DEFINITION

The output guaranty metric is derived by comparing actual and forecasted time series at any selected time horizon and granularity. This metric corresponds to the cost of the smallest possible energy buffer that could make up for all instances of missed forecast over a given test period. The metric is normalized to a nominal size of installed PV capacity equal to 1 kW. The metric requires a definition of the buffer storage's technical specs -- cost per kWh capacity and efficiency -- and a definition of the dispatching algorithm (these definitions have not yet been standardized.) The dispatching algorithm presented here assumes: (1) that excess production with respect to forecast is either curtailed or used to charge the storage, (2) that production below forecast is met via stored energy, and (3) that the buffer can be fully recharged in low demand hours.

The ramp rate mitigation metric requires a determination of the amount of storage sufficient to bring maximum acceptable ramps below a selected threshold by applying a running mean algorithm. Figure 1 at right shows the application of such a running mean algorithm without forecast (Figure 1, B1) and with perfect forecast (Figure 1, B2). The storage specs and dispatch algorithm are similar to those defined above: excess production relative to the running mean is either use to recharge the battery or curtailed, overestimated production is made up by storage, storage is re-initialized in off hours. The forecast accuracy metric per se consist of comparing the supplemental cost of storage incurred when applying the forecast versus an ideal forecast. Forecasting skill can also be determined by comparing this supplemental amount of storage to the amount required with a persistence forecast (Figure 1, B1). As above this metric is normalized to 1 kW PV installed capacity.

III. OPERATIONAL EXAMPLES

We provide examples of metric application for regional fleet forecast evaluation in two climatically distinct US regions centered respectively on the SURFRAD sites of Bondville and Desert Rock. Considered footprints range from a single point to balancing areas the size of California. The forecast models analyzed include SolarAnywhere V4 as well as its underlying models: ECMWF, GFS, HRRR, NDFD and satellite-derived cloud motion.

For output guaranties the considered time horizons include 1, 3, 24 and 48 hours-ahead with an hourly time granularity. For ramp rate reduction, we consider an objective of less than 25% of installed capacity per hour.

For these examples we assume a life-cycle buffer storage cost of \$300/kWh and a round trip efficiency of 90%, corresponding to the expected cost and specs of electrical batteries in the foreseeable future. The results, quantified in terms of operational storage cost, can be applied directly to quantify the cost of injecting variable, partially predictable PV power onto power grids.

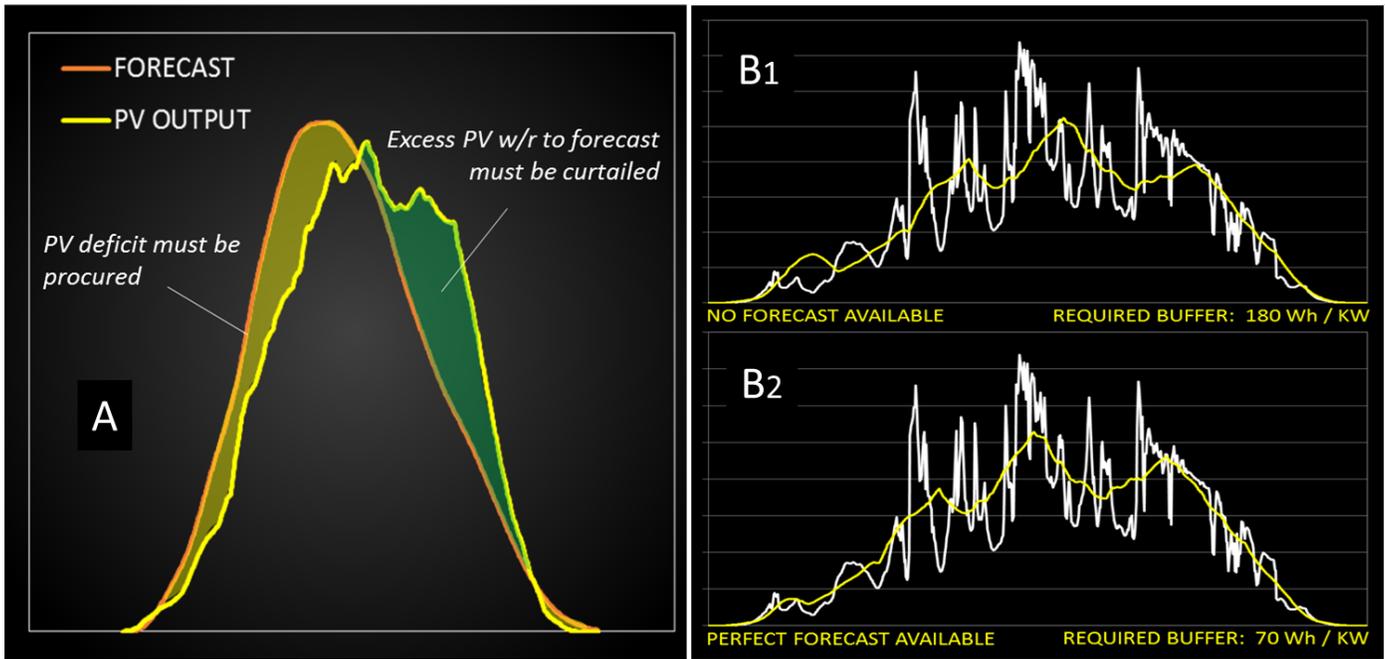


Figure 1 – (A) comparing forecasted and actual PV fleet output. A fleet of storage systems procures the missing energy when the forecast overestimates production and absorbs excess energy otherwise. The minimum possible storage achievable with a forecast model is used as the performance metric for that model. (B) A storage is driven by running mean algorithm to set the maximum acceptable ramp rate below a desired threshold. If future output is not known (B1), the running mean algorithm is based on the system’s history (smart persistence). If it can be forecasted ideally (B2) the running mean algorithm can take advantage of future production, resulting in considerably reduced storage requirements. The storage requirements achieved with a given forecast model in relation to these ideal forecast and no forecast cases amount to the accuracy metric for that model.