ISSN: 2519-9439

Machine Learning-Based Energy Management in Solar-Powered Smart Grids: Integrating Demand Forecasting and Battery Storage Optimization

Sofia Lindholm

Centre for Intelligent Energy Networks, Westbridge International Research Institute, Finland

Abstract

Solar-powered smart grids must coordinate variable photovoltaic generation, changing electricity demand, and constrained storage resources. Machine learning can improve the representation of demand and generation patterns, while optimization can translate forecasts into battery charging and discharging decisions. However, predictive accuracy alone does not establish reliable or economical energy management. This narrative review integrates established literature on probabilistic load forecasting, solar prediction, machine learning, microgrid control, battery degradation, and prescriptive analytics. It proposes a framework linking data provenance, forecast development, uncertainty assessment, physically feasible scheduling, and operational monitoring. Particular attention is given to chronological validation, correlated load and solar errors, state-of-charge limits, conversion losses, battery aging, and the distinction between scheduling and electrical protection. An illustrative graph compares grid imports before and after battery dispatch in a synthetic daily optimization example with explicitly stated assumptions. The calculation uses known profiles and therefore does not demonstrate realized forecasting benefits or field savings. A comparative table summarizes complementary analytical approaches and their main limitations. The analysis argues that probabilistic forecasts, appropriate benchmarks, and realistic closed-loop evaluation are necessary to establish operational value. Storage schedules should reflect reserve requirements and battery condition rather than prioritize short-term electricity costs alone. Future empirical studies should compare complete forecasting and control systems under seasonal variation, communication failures, and changing demand. The proposed framework supports reproducible research and local evaluation of reliable, efficient, and accountable energy management.

Keywords: machine learning, solar smart grids, demand forecasting, battery storage optimization, energy management

1. Introduction

Photovoltaic generation introduces an energy resource whose availability depends on sunlight, weather, installation characteristics, and equipment condition. Electricity demand follows different patterns, shaped by occupancy, industrial activity, temperature, and consumer behavior. These patterns do not necessarily coincide. A site may produce surplus solar electricity around midday while purchasing substantial electricity during the evening. Battery storage can shift energy across time, but its usefulness depends on available capacity, power limits, conversion losses, and the value attached to alternative uses of stored energy.Smart grids connect measurements, communication, control, and distributed energy resources to support coordinated operation. Microgrids represent a related setting in which

ISSN: 2519-9439

local resources and loads are managed together, potentially with both grid-connected and islanded operation. Lasseter (2002) provides an early foundation for the microgrid concept. Solar-powered does not necessarily mean electrically independent: many practical systems continue to import and export electricity. The intended mode of operation should therefore be specified before evaluating an energy management method.Machine learning contributes by estimating future demand or generation from historical measurements and contextual variables. Optimization contributes by selecting actions under objectives and constraints. A forecast that minimizes average error can still lead to expensive or infeasible dispatch if its largest errors occur during critical periods. Conversely, a modest forecast can support useful decisions when uncertainty is represented appropriately and reserves are adequate. The relationship between prediction and action is consequently central to evaluating the full energy management system.This article presents a focused narrative review and a proposed framework for combining demand forecasting with battery optimization. Twelve publications provide methodological foundations covering load and solar forecasting, recurrent and tree-based models, probabilistic scoring, storage technologies, battery aging, microgrid control, and prescriptive analytics. The source selection is thematic rather than an exhaustive systematic search. No machine learning model is trained in this article, and the numerical illustration does not estimate an average improvement across installations or establish the superiority of a specific forecasting architecture.The framework distinguishes three questions. The forecasting question concerns what demand and solar output may occur over the relevant horizon. The scheduling question concerns when storage should charge, discharge, or preserve energy. The operational question concerns whether those schedules remain feasible and useful under actual network conditions. These questions require different evidence. A low forecast error does not demonstrate compliant voltage behavior, and a feasible energy balance does not establish economic value after degradation and implementation costs.The following nine sections develop the framework from system objectives through deployment. They emphasize explicit definitions, chronological evaluation, physically defensible constraints, and continuing feedback. Hong and Fan (2016) explain the importance of probabilistic load forecasts and their evaluation, while Parisio et al. (2014) demonstrate a model predictive control approach to microgrid operation. The argument here is that their complementary perspectives should be integrated: uncertainty about the future must be evaluated in relation to the actions and constraints it informs.

2. System Objectives and Smart Grid Operating Context

Energy management begins with a precise system boundary. A household installation, commercial campus, industrial microgrid, and distribution feeder have different controls and responsibilities. The available assets may include rooftop photovoltaics, stationary batteries, electric vehicles, flexible loads, and a connection to the host grid. The operator should identify which actions can actually be commanded and which quantities are externally imposed. A forecasting study becomes difficult to interpret when its claimed benefits depend on flexibility that the installation does not possess.Objectives may include reducing electricity expenditure, limiting peak imports, increasing solar self-consumption, or preserving service during interruptions. These objectives can conflict. Keeping a battery full for backup can reduce the capacity available to absorb midday generation. Avoiding a demand peak can require discharging energy that would otherwise serve a later critical load. An objective function should therefore state priorities and the terms used to represent them. A combined monetary score requires defensible valuations rather than arbitrary weights chosen to produce attractive results.The control hierarchy also matters. Olivares et al. (2014) distinguish primary, secondary, and tertiary functions in microgrid control. A supervisory schedule

ISSN: 2519-9439

operating over minutes or hours should not be presented as a replacement for fast inverter controls, protection, or safe transition procedures. Grid-connected scheduling and islanded resilience also require different assumptions about power balance and available support. The proposed framework places forecasting and battery optimization within the supervisory layer, supported by lower-level controls that enforce equipment limits and maintain acceptable electrical operation.

3. Metering Weather Data and Forecast Preparation

Demand and solar forecasting depend on measurements with consistent time references and physical meanings. Electrical power is measured in kilowatts, while energy over an interval is measured in kilowatt-hours. Confusing these quantities can create incorrect dispatch equations or misleading performance claims. Data preparation should document sampling intervals, time zones, daylight-saving adjustments where applicable, meter coverage, and missing observations. The same interval boundaries should be used when combining load, photovoltaic output, prices, and battery measurements.Gross demand should be distinguished from net demand at the grid connection. A meter behind a solar installation may record imports after local generation has already served part of the load. Training on this signal and then subtracting solar output again would double-count generation. Similarly, battery actions can change the observed net-load history. Researchers should define the target carefully and retain component-level measurements where feasible. If gross load must be reconstructed, uncertainty and assumptions in that reconstruction should accompany the forecasts.Weather inputs require attention to what was available when the schedule was produced. Observed future irradiance cannot be treated as an operational forecast. Historical weather forecasts, or a clearly justified substitute, should be used in realistic backtesting. Features may include prior demand, calendar effects, temperature forecasts, cloud information, and recent solar output. Imputation, scaling, and feature selection should be fitted within development periods. These practices help prevent information leakage and preserve comparability between the analytical dataset and the system's actual information conditions.

4. Machine Learning for Electricity Demand Forecasting

Demand forecasting should start with credible reference methods. Seasonal persistence, recent averages, and calendar-based regression can be strong baselines when patterns repeat. Their performance establishes whether a complex model provides useful additional information. The comparison should use identical forecast origins, horizons, covariates, and validation periods. A model intended for hourly battery scheduling should be tested on the horizons required by that scheduler rather than only on an easier one-step prediction task.Tree-based boosting can represent nonlinear relationships between demand, temperature, recent history, and calendar variables. XGBoost provides a scalable boosting framework that can be adapted to such tabular forecasting tasks (Chen and Guestrin, 2016). Recurrent neural networks provide another option for sequential information; long short-term memory architectures were introduced by Hochreiter and Schmidhuber (1997). Neither method guarantees better energy management. Their value depends on dataset size, regime stability, feature construction, and the practical costs of model training and maintenance.Forecast development should account for multiple horizons and changing behavior. Direct horizon-specific models can differ from recursive approaches that reuse previous predictions. Electric vehicle charging, new equipment, occupancy changes, and demand response can alter patterns that were reliable during training. Global models can share information across meters, but aggregation and privacy arrangements require attention.

ISSN: 2519-9439

Demand forecasts should be assessed during high-load periods as well as average conditions. Errors near critical import peaks may matter more for storage decisions than equally sized errors during low-cost, unconstrained intervals.

5. Solar Forecasting and Joint Uncertainty Representation

Solar forecasts may target irradiance or delivered photovoltaic power. These are related but distinct quantities because conversion depends on module temperature, inverter behavior, orientation, shading, and equipment availability. Clear-sky patterns provide useful structure, while clouds introduce short-term variation that can be difficult to anticipate. Voyant et al. (2017) review machine learning methods for solar radiation forecasting. Applying these methods to operational power prediction requires a defensible link between environmental inputs and the installation's measured electrical output.Uncertainty in solar generation should be represented alongside uncertainty in demand. Independent marginal intervals can misrepresent net demand when load and generation errors are related through weather or other common influences. Scheduling may require scenarios that preserve these relationships across variables and time. An isolated low-solar prediction interval does not specify the likelihood of a prolonged cloudy sequence coinciding with high demand. The representation should match the planning horizon and the events that could exhaust storage or exceed import limits.Probabilistic forecasts require evaluation of both calibration and useful concentration. Very wide intervals can achieve coverage without supporting precise actions, while narrow but unreliable intervals can encourage excessive risk. Gneiting and Raftery (2007) establish principles for scoring predictive distributions. Hong and Fan (2016) provide energy-specific context for probabilistic demand forecasting. Scenario evaluation should include tails and operationally important periods. A visually plausible fan chart is insufficient evidence that a forecast distribution accurately represents the risks relevant to battery scheduling or reserve decisions.

6. Battery Dynamics Efficiency and Aging Constraints

A storage model should distinguish energy capacity from power capability. Capacity limits the quantity of stored energy; charging and discharging power limits determine how quickly it can change. For a time interval, the next stored-energy state equals the current state plus charging energy multiplied by charging efficiency, minus discharged energy divided by discharging efficiency. State-of-charge bounds, converter ratings, and measurement uncertainty restrict feasible operation. Vazquez et al. (2010) review storage technologies and applications, showing why device characteristics must be matched to their intended service.Optimization should also avoid nonphysical or commercially meaningless schedules. Charging and discharging simultaneously may need explicit prevention, depending on the formulation, prices, and loss assumptions. Import and export constraints, inverter coupling, and unavailable equipment should be represented consistently. A terminal stored-energy requirement prevents the scheduler from appearing economical simply by emptying the battery at the end of the evaluation horizon. Reserve constraints may be necessary when the same battery also supports critical loads or interruption preparedness.Battery aging changes the long-term value of dispatch. Vetter et al. (2005) review lithium-ion aging mechanisms, while Xu et al. (2018) model degradation for life assessment using operating profiles. Cycle depth, temperature, time, and operating conditions can affect degradation, with relationships depending on chemistry and design. A simple throughput penalty can discourage unnecessary cycling but is not a full aging model. Reports should state the approximation and test its

ISSN: 2519-9439

influence. Ignoring degradation can favor short-term arbitrage that is unattractive when replacement and reduced usable capacity are considered.

7. Optimization and Receding Horizon Energy Management

A battery scheduler translates forecasts into charging, discharging, import, export, and possibly curtailment decisions. A basic objective can minimize energy purchase costs net of export revenue, with additional terms for degradation or unmet demand where relevant. Constraints enforce interval power balance, energy-state evolution, and equipment limits. More detailed settings may require mixed-integer formulations for discrete equipment states or network constraints. The level of detail should reflect the question rather than imply that an aggregate energy model represents every electrical phenomenon.Model predictive control repeatedly solves a finite-horizon scheduling problem using updated measurements and forecasts. Only the initial action or short segment is implemented before the problem is solved again. Parisio et al. (2014) apply this perspective to microgrid operation. Reoptimization can respond to forecast errors and changing conditions, but its performance depends on model fidelity, computational timing, and fallback behavior. A schedule calculated once with complete knowledge of the next day is an idealized planning calculation, not a demonstration of closed-loop control effectiveness.Stochastic scheduling can optimize over plausible scenarios, while robust methods protect against specified uncertainty sets. Both require explicit assumptions about risk and feasibility. Excessive conservatism may preserve reserves at high cost; insufficient protection can leave the battery unable to meet an important obligation. Bertsimas and Kallus (2020) connect contextual prediction to prescriptive decisions, supporting evaluation of the actions induced by forecasts. The proposed framework therefore compares complete forecast-and-control combinations rather than selecting a predictor only by its average error.

Table 1: Complementary approaches for solar and battery energy management.

Approach	Useful contribution	Main limitation
Seasonal baseline	Transparent demand reference	Limited adaptation to regime change
Boosting model	Nonlinear use of demand covariates	Sensitive to feature timing and tuning
Sequence model	Represents temporal dependencies	Data and maintenance requirements
Probabilistic forecast	Quantifies future uncertainty	Needs calibration and joint scenarios
Battery optimization	Schedules energy within constraints	Depends on costs and physical fidelity
Receding-horizon control	Updates actions with new information	Requires timely solves and fallback

ISSN: 2519-9439

8. Demand Flexibility Network Limits and Resilience

Flexible demand can complement battery storage when users can shift selected activities without unacceptable inconvenience or production loss. Examples include charging schedules and some heating or industrial processes, subject to their own constraints. A shifted load is not automatically eliminated energy use, and delayed consumption can create a later rebound. Optimization should represent deadlines, comfort or process requirements, and realistic participation. Assuming unlimited flexibility can exaggerate benefits and obscure the contribution of the battery itself.Aggregate power balance does not guarantee acceptable operation across a distribution network. Voltage limits, branch loading, phase imbalance, and inverter reactive-power capability can constrain dispatch. A site-level model may reasonably omit these details for an initial economic analysis, but should identify that boundary. If the research claims network-wide feasibility, suitable electrical modeling and validation are needed. Olivares et al. (2014) provide context for coordinating microgrid control functions, while Lasseter (2002) highlights the broader organization of local generation and loads.Resilience requires separate evaluation from routine electricity-cost reduction. Islanded operation depends on available generation, storage, grid-forming capability, and load prioritization. A battery that reduces a daily bill may not preserve sufficient energy for an interruption. Reserve policies should specify critical demand, assumed interruption duration, and acceptable unmet energy. Stress tests can examine prolonged low-solar conditions and communication failures, but their conclusions remain conditional on the selected scenarios. Lower imports or greater solar self-consumption should not be presented as direct proof of improved resilience.

9. Validation Benchmarks and Operational Performance

Forecast evaluation should use chronological partitions and rolling origins that reproduce the information available to the scheduler. Randomly mixing adjacent intervals can leak seasonal or behavioral information. Tuning should remain separate from final evaluation, and all competing models should use comparable inputs. Mean absolute and squared errors offer complementary information, but results should be separated by horizon, season, and operating regime. Probabilistic evaluation should assess coverage, quantile performance, and calibration during periods when storage decisions are most consequential.Controller evaluation requires stronger comparisons than a single attractive daily profile. Baselines can include no storage, a simple self-consumption rule, and an established tariff-based scheduler. When assessing the value of machine learning specifically, storage hardware and control objectives should be held constant while forecasting methods change. A perfect-information schedule can provide a useful comparison bound under stated assumptions, but it should not be confused with achievable operation. Repeated seasonal testing is needed because a sunny day can favor a policy that performs poorly during winter or variable cloud conditions.Operational measures should match the stated objectives. Relevant outcomes include electricity expenditure, peak imports, curtailed energy, battery throughput, constraint violations, reserve availability, and unmet critical demand. Self-consumption and self-sufficiency should be defined explicitly because their denominators differ. Statistical uncertainty and sensitivity to tariffs, battery size, efficiency, and degradation assumptions should be reported. Lower energy bills do not automatically establish lower emissions, which depend on the generation displaced and the timing of charging. The full comparison should account for implementation and storage costs when making investment claims.

ISSN: 2519-9439

10. Deployment Governance and Continuing Adaptation

The proposed framework connects system definition, reliable data, benchmarked forecasts, uncertainty assessment, constrained scheduling, and operational monitoring. Initial deployment can use shadow operation to check forecasts and proposed schedules without issuing commands. Hardware-in-the-loop or comparable testing can then examine timing and equipment interfaces before a controlled pilot. These stages answer different questions. Successful offline optimization does not establish reliable communications, correct device responses, or safe behavior during an unexpected loss of information.Operational governance should define who can alter models, tariffs, reserve requirements, and battery limits. Version control should preserve the configuration associated with each schedule. Meter and battery telemetry require access controls and integrity checks because incorrect inputs can produce inappropriate actions. Customer-level demand data can reveal occupancy and activity patterns, so collection and retention should be proportionate to the task. These are general governance recommendations rather than a claim that a specific architecture satisfies every applicable regulatory or contractual requirement.Continuing evaluation should monitor forecast drift, missing inputs, command execution, energy-state errors, constraint margins, and changes in usable battery capacity. A fallback controller should remain available when the optimizer times out or forecast quality deteriorates. Retraining needs renewed validation of the complete system, particularly when demand patterns or inverter settings change. Expansion to additional sites should follow evidence of transferability rather than assumption. Accountable energy management requires that operators can understand the basis of schedules, identify limitations, and suspend automation when operational evidence no longer supports its use.

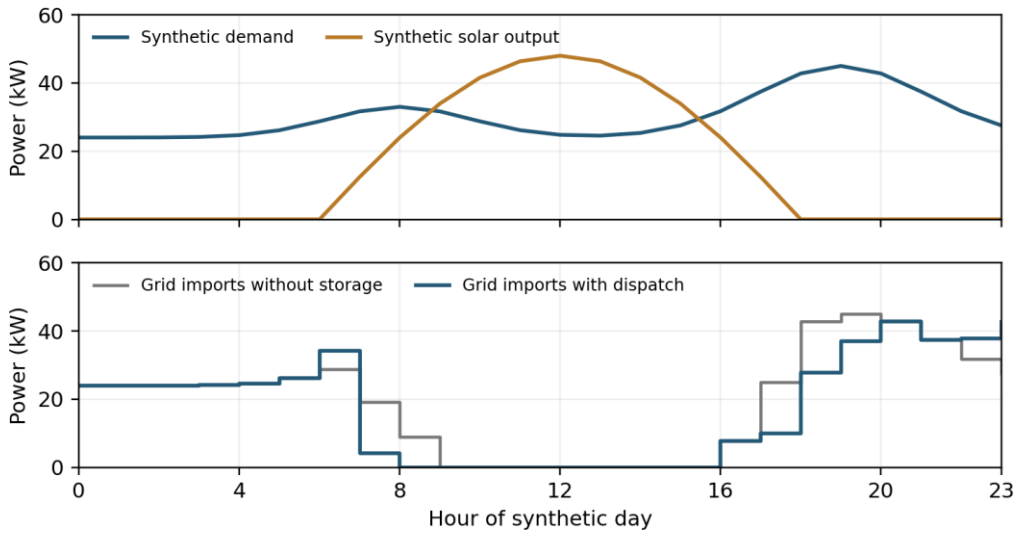


Figure 1: Synthetic demand, solar output, and calculated grid imports with and without battery dispatch. Known-profile illustration; no field or trained-model results.

13. Summary

Machine learning can support solar-powered smart grids by anticipating electricity demand and photovoltaic availability, while battery optimization converts those estimates into actions. Reliable development begins with a defined operating boundary, clear objectives, and measurements that distinguish gross demand, generation, storage activity, and grid exchange. Forecasts should be benchmarked chronologically and evaluated across relevant horizons and

ISSN: 2519-9439

conditions. Complex models can represent useful relationships, but their practical value depends on the decisions they support rather than predictive sophistication alone.Battery scheduling requires physically consistent energy dynamics, power limits, conversion efficiencies, reserve policies, and defensible treatment of degradation. Probabilistic forecasts can represent uncertainty, but calibration and joint scenario behavior require assessment. Receding-horizon control can adapt decisions as information changes, provided equipment interfaces and fallback behavior are dependable. Supervisory scheduling remains distinct from fast electrical control and protection. The synthetic example demonstrates how constrained dispatch can shift grid purchases under known profiles; its numerical differences are not evidence of real-world savings or trained-model effectiveness.The proposed framework links forecasting, optimization, operational testing, and continuing oversight. Evaluations should compare complete systems against meaningful baselines while separating cost, resilience, environmental, and investment claims. Changes in demand, battery condition, and network operation require renewed assessment. Future research should test integrated methods over multiple seasons and sites, including forecast errors, communication failures, and realistic reserve requirements. The main implication is that energy management creates value when reliable predictions lead to feasible actions under actual constraints, with transparent assumptions and evidence that supports continued operation.

References

[1] Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. Management Science, 66(3), 1025–1044. https://doi.org/10.1287/mnsc.2018.3253

[2] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794. https://doi.org/10.1145/2939672.2939785

[3] Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359–378. https://doi.org/10.1198/016214506000001437

[4] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[5] Hong, T., & Fan, S. (2016). Probabilistic electric load forecasting: A tutorial review. International Journal of Forecasting, 32(3), 914–938. https://doi.org/10.1016/j.ijforecast.2015.11.011

[6] Lasseter, R. H. (2002). MicroGrids. IEEE Power Engineering Society Winter Meeting, 1, 305–308. https://doi.org/10.1109/PESW.2002.985003

[7] Olivares, D. E., et al. (2014). Trends in microgrid control. IEEE Transactions on Smart Grid, 5(4), 1905–1919. https://doi.org/10.1109/TSG.2013.2295514

[8] Parisio, A., Rikos, E., & Glielmo, L. (2014). A model predictive control approach to microgrid operation optimization. IEEE Transactions on Control Systems Technology, 22(5), 1813–1827. https://doi.org/10.1109/TCST.2013.2295737

[9] Vazquez, S., Lukic, S. M., Galvan, E., Franquelo, L. G., & Carrasco, J. M. (2010). Energy storage systems for transport and grid applications. IEEE Transactions on Industrial Electronics, 57(12), 3881–3895. https://doi.org/10.1109/TIE.2010.2076414

[10] Vetter, J., et al. (2005). Ageing mechanisms in lithium-ion batteries. Journal of Power Sources, 147(1–2), 269–281. https://doi.org/10.1016/j.jpowsour.2005.01.006

[11] Voyant, C., Notton, G., Kalogirou, S., Nivet, M.-L., Paoli, C., Motte, F., & Fouilloy, A. (2017). Machine learning methods for solar radiation forecasting: A review. Renewable Energy, 105, 569–582. https://doi.org/10.1016/j.renene.2016.12.095

ISSN: 2519-9439

[12] Xu, B., Oudalov, A., Ulbig, A., Andersson, G., & Kirschen, D. S. (2018). Modeling of lithium-ion battery degradation for cell life assessment. IEEE Transactions on Smart Grid, 9(2), 1131–1140. https://doi.org/10.1109/TSG.2016.2578950