A Proposed Forecast–Coordinate–Constrain Framework for Hydro–PV–BESS Energy Management in Smart Sustainable Communities

Oihika Arpit

Graduate Researcher, ECE, upGrad, India

Jasreet Kaur

Undergraduate Researcher, Bachelor of Science in Analytics, Golden Gate University, USA

Dr. Smrite Goudhaman

Global Professor of Practice, Golden Gate University, USA

Manuscript prepared from the Jabalpur Hydro–PV–BESS simulation and ITU-T smart-community contribution.

Abstract

The rapid addition of variable renewable generation is changing the role of energy management from simple source prioritization to anticipatory coordination. This paper presents a reproducible research-design framework for predictive energy management of a hydro–photovoltaic–battery energy storage system (Hydro–PV–BESS), using a Jabalpur-based simulation as the engineering case study. The existing Python prototype integrates a 55 MW solar photovoltaic (PV) resource, a 70 MWh battery energy storage system (BESS), hydropower reserve inspired by the Rani Avanti Bai Sagar Hydel Power Station, and fossil-grid fallback. Its documented 24-hour baseline supplies 82.98% of demand from renewable sources, while 180.57 MWh remains dependent on fossil fallback. The baseline uses a transparent reactive sequence in which PV serves demand first, followed by BESS discharge, hydropower, and finally fossil supply. The present paper does not claim that an artificial intelligence (AI) controller has already improved these results. Instead, it formalizes a Forecast–Coordinate–Constrain architecture in which short-horizon forecasts are converted into a constrained rolling-horizon dispatch problem. The framework specifies power-balance,

storage, reserve, curtailment, and reliability constraints required for implementation and defines a validation protocol covering forecast error, renewable supply share, fossil dependence, battery throughput, terminal state of charge (SOC), curtailment, reliability, and emissions. It further extends the original one-day demonstration into a research design for multi-day and seasonal testing, sensitivity analysis, and degradation-aware operation. The contribution is therefore an engineering integration and reproducibility framework, not a new forecasting algorithm or a demonstrated AI performance claim. The design also maps its operational data and key performance indicators (KPIs) to the emerging International Telecommunication Union Telecommunication Standardization Sector (ITU-T) smart-city energy-management and residential energy-storage recommendations.
Keywords: Hydro–PV–BESS; predictive energy management; renewable forecasting; constrained optimization; smart sustainable communities.

1. Introduction

Electricity systems are becoming less predictable as the contribution of variable renewable energy grows. Solar photovoltaic (PV) generation changes with weather and daylight, while electricity demand retains strong intraday and seasonal structure. A hybrid portfolio that combines PV, battery storage, and hydropower can soften these variations, but only when the complementary strengths of the three resources are coordinated rather than operated as independent priority blocks. Recent energy-management literature distinguishes rule-based control, mathematical optimization, model predictive control, and learning-based approaches precisely because the operational problem is no longer only to decide what resource is available; it is to decide what resource should be preserved, released, or reserved for a future condition (Zhu et al., 2025; Pengtem et al., 2025; Zambrano et al., 2025; Phillips et al., 2026).

This challenge is particularly relevant in Madhya Pradesh. The Central Electricity Authority (CEA) projects a 5.36% compound annual growth rate in electricity demand through 2034–35 and notes pronounced seasonal and diurnal variation in the state demand profile. The same assessment identifies the need for methodical resource planning and emphasizes the growing importance of solar generation and battery storage in the future generation mix (Central Electricity Authority, 2024). For a smart-community energy platform, that trajectory suggests that dispatch logic should be evaluated not only by how much renewable electricity is available today, but by how effectively the system preserves flexibility for the next several hours.

The present study begins from a deliberately simple and transparent engineering artifact: a 24-hour Python simulation of a hybrid Jabalpur energy system. The repository documents a dispatch sequence of solar PV → battery storage → hydropower reserve → fossil-grid fallback, together with battery state-of-charge tracking, hourly demand-supply balancing, and carbon-footprint estimation. The published baseline reports 1,061.00 MWh of daily demand, 347.31 MWh of solar delivered directly to load, 33.24 MWh of solar used to charge the battery, 53.28 MWh delivered by the battery, 479.84 MWh dispatched by hydropower, and 180.57 MWh of fossil fallback (Arpit, 2026). The model is explicitly described as an offline educational simulation rather than a connection to real plant supervisory control and data acquisition (SCADA) or control equipment (Arpit, 2026).

The research problem is therefore narrower and more useful than the claim that AI will automatically improve renewable-energy performance. The question is: how can a reproducible rule-based hybrid simulation be transformed into a predictive, constraint-aware energy management architecture without losing the transparency that makes the original model useful? This paper addresses that question by formulating a Forecast–Coordinate–Constrain framework in which forecasting supplies anticipatory information, coordination assigns roles to PV, BESS, and hydro, and constrained optimization prevents the predictive controller from proposing physically infeasible schedules.

1.1 Research objectives

- Establish a transparent baseline for hydro–PV–BESS dispatch using the Jabalpur simulation.
- Identify the operational weakness of purely reactive dispatch under forecastable demand and renewable variability.
- Formulate an AI-enabled predictive architecture that combines short-horizon forecasting with constraint-aware coordination.

• Define measurable performance indicators that allow the future predictive controller to be compared fairly with the existing rule-based baseline.
• Connect the technical design to smart-sustainable-community requirements for energy data, storage services, and Sustainable Development Goal-oriented performance measurement.

2. Case Study and Baseline System

The case study is centered on Jabalpur, Madhya Pradesh, and uses hydropower reserve characteristics inspired by field exposure to the Rani Avanti Bai Sagar Hydel Power Station at Bargi Nagar. Public descriptions of the station report two 45 MW generating units, while the present simulation abstracts the hydro resource into a renewable reserve rather than reproducing plant hydraulics or operational control (Madhya Pradesh Power Generating Company Limited, 2026; Arpit, 2026). This distinction is important: the study uses the plant as an engineering context for the hybrid model, not as a source of confidential operational data.
The hydropower context is informed by Oihika Arpit's technical training exposure at the Rani Avanti Bai Sagar Hydel Power Station, Bargi Nagar, Jabalpur. The exposure included observation of hydropower-generation equipment and operating systems relevant to the engineering interpretation of the hydro component. This field exposure informed the case-study context only; no confidential plant operational data are used in the simulation.

Table 1.

Baseline system configuration used in the Jabalpur simulation.

Component	Modelled value / role	Study interpretation
Solar PV	55 MW	Primary daytime renewable source
BESS	70 MWh	Stores excess PV and supports deficit periods
Hydropower reserve	Hydro reserve; inspired by 2 × 45 MW Bargi station	Dispatchable renewable balancing resource

Component	Modelled value / role	Study interpretation
Fossil fallback	180.57 MWh in baseline	Residual supply when renewable resources are insufficient
Simulation horizon	24 h	Initial reproducible daily operating case

The original Python configuration and executable logic were reviewed for the baseline parameters. The code explicitly defines the 24-hour horizon (Δt = 1 h), 55 MW synthetic PV capacity, 70 MWh BESS, initial SOC of 45%, SOC limits of 10–90%, charge and discharge efficiencies of 0.95, 16 MW charge/discharge power limits, a 34.10 MW hydro reference capacity, a synthetic hourly PV profile with fixed cloud factors, a synthetic 24-hour demand array, explicit PV curtailment logic, and an emission factor of 820 kg CO_2/MWh. These values are therefore documented as model parameters rather than left as unrecoverable assumptions (Arpit, 2026).

Table 2.

Consolidated modelling assumptions and documentation status.

Parameter	Current baseline	Status for publication
Simulation horizon	24 h	Measured/modelled baseline; longer validation required
Solar PV capacity	55 MW	Documented simulation configuration
BESS energy capacity	70 MWh	Documented simulation configuration
Daily demand	1,061.00 MWh	Documented simulation result/input profile
BESS charge/discharge efficiency	η_{ch} = 0.95; η_{dis} = 0.95	Explicitly defined in configuration and used by the battery model
Initial SOC	45% = 31.50 MWh	initial_soc_pct = 45.0
SOC limits	10%–90% = 7–63 MWh	min_soc_pct = 10.0; max_soc_pct = 90.0
PV profile source	Synthetic code-generated profile	solar_profile(): daylight sine function × fixed 24-hour cloud-factor vector

Parameter	Current baseline	Status for publication
Demand profile source	Synthetic/fictitious 24-hour regional demand	Explicit DEMAND_MW array in code; not measured Jabalpur demand
Hydro availability	34.10 MW × fixed hourly availability factor	hydro_available_profile() explicitly defines the hourly availability-factor vector
Hydro ramping/reservoir	Excluded	Must be added for higher-fidelity study
Fossil fallback	Residual supply	Used after renewable resources are insufficient
Curtailment rule	max(0, solar_surplus − battery_charge)	Explicitly implemented after battery charging
Emission factor	820.0 kg CO_2/MWh	fossil_emission_factor_kg_per_mwh = 820.0; source should be cited/archived
Control status	Offline simulation	No live SCADA/programmable logic controller (PLC)/turbine/inverter/battery control

The baseline is reproducible from the supplied Python configuration: the battery starts at 45% SOC, operates between 10% and 90% SOC, uses 95% charge and discharge efficiency, and is limited to 16 MW charging/discharging. PV is generated synthetically from a daylight sine function multiplied by a fixed 24-hour cloud-factor vector; demand is a synthetic 24-hour array explicitly labelled as fictional regional demand; hydro availability is generated from a 34.10 MW reference capacity multiplied by a fixed hourly availability-factor vector; excess PV is curtailed after the battery charging limit is reached; and fossil carbon is calculated using 820 kg CO_2/MWh (Arpit, 2026). Because the PV variable represents instantaneous power or installed generation capacity, it is correctly expressed in MW; the BESS energy capacity is expressed in MWh. This distinction is retained consistently throughout the manuscript to avoid conflating power and energy units.

2.2 Baseline dispatch logic

The repository implements a deterministic dispatch order. Solar generation serves the instantaneous load first; excess solar is directed to the battery; when renewable generation is insufficient, the battery is discharged; hydropower is then dispatched as renewable backup; and fossil-grid supply covers the remaining deficit. The baseline code also explicitly computes curtailed solar as the portion of solar surplus remaining after the battery charging action. The public repository identifies the simulation as an offline educational model (Arpit, 2026).

Its weakness is the same property that makes it simple: it is primarily present-state driven. A present-state decision can be locally reasonable while being globally poor over a short horizon. For example, discharging the battery during one deficit hour may reduce fossil use immediately but leave too little stored energy for a larger evening deficit. Similarly, using hydropower whenever the battery is empty can reduce current fossil dependence but consume renewable flexibility that might be more valuable later. A predictive energy management system (EMS) should therefore evaluate resource actions against anticipated future conditions rather than only the current power mismatch.

Hydropower is responsible for a large portion of the renewable balancing in the baseline, so its abstraction is a first-order modelling issue. The current prototype represents hydro as an available renewable reserve but does not reproduce reservoir mass balance, turbine efficiency curves, hydraulic head variation, minimum releases, unit commitment, ramp-rate limits, start-up constraints, water opportunity cost, or seasonal hydrological uncertainty. These omissions can make the hydro block appear more flexible than an actual plant. Accordingly, the reported 479.84 MWh of hydro dispatch should be interpreted as a simulation result under the stated reserve abstraction rather than as a prediction of actual Rani Avanti Bai Sagar plant operation. The next experimental version should introduce time-varying hydro availability and ramping constraints at minimum; a high-fidelity version should additionally model reservoir state and water value.

2.3 Reported baseline results

The following results are taken from the documented GitHub prototype and are presented as baseline simulation evidence, not as results of a completed AI experiment (Arpit, 2026).

Table 3.

Metrics reported by the GitHub prototype.

Metric	Baseline result	Interpretation
Daily load	1,061.00 MWh	24-hour simulated demand
Solar direct to load	347.31 MWh	Direct PV contribution
Solar to battery	33.24 MWh	Energy shifted through storage
Battery delivered	53.28 MWh	Discharge contribution
Hydropower dispatched	479.84 MWh	Renewable balancing contribution
Fossil fallback	180.57 MWh	Residual non-renewable dependence
Renewable supply	82.98%	Baseline renewable share
Final battery SOC	7.00 MWh (10%)	Low terminal storage state
Hydro dispatch hours	18 h	Hydro was used in most hours
Fossil fallback hours	8 h	Deficit hours requiring fossil supply
Carbon footprint	148,063.60 kg CO_2	Reported baseline estimate

Two results are particularly informative for the next research stage. First, the renewable contribution is already substantial at 82.98%, demonstrating that coordinated resource diversity can materially reduce fossil dependence even with simple control. Second, the 180.57 MWh fossil requirement and the final 10% battery SOC reveal the limits of the baseline under the selected assumptions. These are not, by themselves, evidence that an AI controller will reduce the deficit. Instead, they define a concrete baseline against which the predictive framework can be tested.

3. Related Work and Research Gap

Energy-management research increasingly uses optimization and learning to handle the temporal coupling of renewable generation, demand, and storage. A recent review of

grid-connected utility-scale renewable hybrid power plants groups EMS approaches into rule-based control, mathematical optimization, model predictive control, deep reinforcement learning, and stochastic dynamic programming (Zhu et al., 2025). This taxonomy is useful for positioning the present work: the Jabalpur prototype belongs to the rule-based category, while the proposed extension moves toward forecast-assisted constrained optimization, leaving room for a later learning-based correction layer.
Recent work on hybrid renewable systems also shows that the operational problem cannot be reduced to source availability alone. Studies combine PV and hydropower with battery storage to improve renewable accommodation and grid performance, while newer research highlights the need to represent hydropower safety limits, intertemporal decisions, and battery degradation (Pengtem et al., 2025; Grimaldi et al., 2024; Tang et al., 2025). Forecast-aware rolling-horizon optimization has likewise been shown to capture an operational advantage over perfect-foresight assumptions because actual scheduling decisions are made from imperfect information (Phillips et al., 2026).
On the forecasting side, machine learning is widely used because short-term PV production is strongly nonlinear and sensitive to weather and operating context. Recent reviews identify neural, deep-learning, and hybrid methods as important options, while also emphasizing generalization and the persistent gap between forecasting research and integrated storage control (Gupta et al., 2025; Casazola-Cruz et al., 2025). The literature therefore supports a modular design in which forecast accuracy is treated as an input to energy management, not as an isolated predictive benchmark.

3.1 Specific gap addressed by this paper

The gap addressed here is architectural. Much of the literature optimizes either the renewable system or the forecasting problem, whereas the present study starts from an inspectable hybrid dispatch prototype and asks how predictive intelligence can be embedded without obscuring physical constraints. The proposed contribution is consequently not a new forecasting algorithm. It is a framework for translating forecast information into physically admissible, resource-aware dispatch decisions for a smart-community energy platform.

4. Proposed Forecast–Coordinate–Constrain Framework

The framework has three operational layers. The first forecasts the near-future load and renewable availability. The second converts those forecasts into a coordinated power requirement across PV, BESS, and hydro. The third enforces physical and operational constraints before the next control action is released. The loop then repeats as new measurements arrive. This architecture retains the original simulation's interpretability while allowing progressively more sophisticated AI components to be introduced.

4.1 Layer 1: short-horizon forecasting

The forecasting layer receives historical load, historical PV output and, where available, weather or irradiance variables. A practical implementation can compare a lightweight tree-based model with sequence models such as long short-term memory (LSTM) or gated recurrent unit (GRU) models and select the model using out-of-sample error rather than model popularity. The paper intentionally does not claim a trained model or forecast accuracy because the current repository documentation does not contain a validated AI forecast experiment. That restraint is important for scientific reproducibility. For load forecasting, the target is the next H hours of demand. For PV, the target is the next H hours of available generation. Hydropower is treated differently because it is dispatchable but bounded; the forecasting layer therefore estimates availability or reserve conditions rather than assuming unlimited hydro output.

4.2 Layer 2: coordinated resource allocation

The coordinator receives forecast trajectories and current states. Its central question is not which source is highest in a static priority list, but which combination of source actions minimizes future dependence on undesirable or constrained resources while preserving feasibility. A useful policy is to value solar first because of its zero marginal fuel requirement, use BESS strategically for temporal shifting and short-term balancing, preserve hydro for deficits where its longer duration or dispatchability is valuable, and treat fossil energy as the residual resource of last resort.

4.3 Layer 3: constrained optimization

At each time step t, the dispatch should satisfy the power balance:

$$P_{load,t} = P_{PV,t} + P_{BESS,t} + P_{H,t} + P_{F,t} \quad (1)$$

$$SOC_{t+1} = SOC_t + \eta_{ch} P_{ch,t} \Delta t / E_B - P_{dis,t} \Delta t / (\eta_{dis} E_B) \quad (2)$$

$$P_{BESS,t} = P_{dis,t} - P_{ch,t} \quad (3)$$

$$0 \le P_{H,t} \le P_{H,max}(t) \quad (4)$$

$$0 \le P_{F,t} \quad (5)$$

where $P_{BESS,t}$ is positive for discharge and $P_{F,t}$ represents fossil fallback. The storage state equation is subject to $SOC_{min} \le SOC_t \le SOC_{max}$ and mutually exclusive charging and discharging. Hydropower is bounded by its allowable power range and reserve policy, while fossil fallback is non-negative. PV output can be curtailed when all compatible sinks are saturated.

The short-horizon objective can then be expressed as a weighted cost function:

$$J = \Sigma_t [w_F P_{F,t} + w_C P_{curt,t} + w_S \Delta SOC_t^2 + w_R R_t + w_U U_t] \quad (6)$$

where P_{curt} is curtailed PV, ΔSOC captures aggressive battery movement, R represents reserve-policy penalties, and U captures any unmet-load slack. The weights are application-specific and should be reported in the final empirical study. A rolling-horizon controller solves this problem over H future intervals but implements only the first action before updating the forecast and system state.

4.4 Where the AI enters the controller

AI is most defensible in this architecture where uncertainty is high and the optimization problem can remain physically constrained. Forecasting is therefore the first insertion point. A second insertion point is residual decision correction, where a learning agent could learn systematic forecast-to-dispatch biases after a rule-based or optimization-based controller has been validated. The hard constraints should remain outside the learning model so that prediction error cannot authorize impossible battery SOC, hydro output, or power-balance states.

4.5 Framework logic and data flow

The workflow below converts the conceptual architecture into a reproducible sequence. The table is presented after this explanatory paragraph to avoid beginning the subsection with a table.

Table 4.

Proposed Forecast–Coordinate–Constrain workflow.

Stage	Inputs	Core operation	Output
Measure	Load, PV, BESS SOC, hydro status	State estimation and data validation	Current system state
Forecast	Historical energy + weather features	Short-horizon AI prediction	P_load,t:t+H and P_PV,t:t+H
Coordinate	Forecasts + resource limits	Rank/allocate renewable flexibility	Candidate schedules
Constrain	Candidate schedules + physical limits	Optimization / feasibility check	Admissible dispatch
Execute & learn	Observed outcome	Apply first control action and update errors	Receding-horizon next state

5. Validation Strategy

A central contribution of this paper is to define how the predictive extension should be evaluated. Comparing only renewable percentage can be misleading because a controller could increase renewable use by cycling the battery excessively, exhausting hydro reserve, or creating larger forecast-risk penalties. Validation must therefore combine energy, storage, reliability, and sustainability indicators.

5.1 Baseline-versus-predictive comparison

- Rule-based dispatch: the existing repository logic, used without modification as the reference controller.
- Predictive EMS without learning-based correction: forecast + constrained optimization.

• Predictive EMS with residual learning: the optimized controller plus a bounded learning-based adjustment layer, activated only after feasibility checks.

5.2 Recommended performance indicators

Table 5.

Suggested validation indicators for the next experimental phase.

Indicator	Definition	Why it matters
Renewable supply share	Renewable energy served / total demand	Primary decarbonization measure
Fossil dependence	Fossil fallback energy / total demand	Measures residual non-renewable reliance
PV curtailment	Curtailed PV / available PV	Shows whether storage/dispatch uses renewable availability
Battery throughput	Charge + discharge energy	Captures cycling burden
Terminal SOC	SOC at end of horizon	Indicates whether current gains consume future flexibility
Forecast MAE/RMSE	Error of load and PV forecasts	Separates predictive quality from control quality
Unmet load	Residual demand not supplied	Reliability constraint
CO_2 footprint	Fallback energy × emission factor	Environmental outcome

The 24-hour case remains useful as a transparent demonstration of the dispatch logic, but it is insufficient for generalized conclusions. The study therefore defines the following validation matrix for the next experimental stage.

Table 6.

Proposed multi-scenario validation matrix.

Scenario	Stress condition	Primary evaluation
High-solar day	High PV / moderate demand	PV curtailment and BESS charging
Low-solar day	Reduced PV availability	Fossil dependence and hydro reserve

Scenario	Stress condition	Primary evaluation
High-evening-peak day	Strong late-day demand	Terminal SOC and reserve preservation
Weekday profile	Typical working-day demand	Normal operating performance
Weekend profile	Shifted demand timing	Forecast robustness
High-hydro availability	Wet-period reserve	Hydro–BESS coordination
Low-hydro availability	Dry/maintenance condition	Reliability under constrained hydro
Forecast-error stress	Biased/noisy forecasts	Predictive-controller robustness

The next experimental study should vary initial battery SOC, BESS energy capacity, hydro availability, forecast error, and evening-peak magnitude. For the objective J, the weights wF, wC, wS, wR, and wU should not be chosen arbitrarily. They should first be normalized to comparable marginal scales, then selected using either a documented preference hierarchy or a Pareto-based multi-objective procedure. A one-at-a-time sensitivity analysis should quantify the effect of each weight on fossil energy, curtailment, battery throughput, and reliability. An ablation study should compare: (i) rule-based dispatch; (ii) forecast-only scheduling; (iii) forecast plus constrained optimization; and (iv) forecast plus constrained optimization with a bounded learning correction. This design ensures that any future improvement can be attributed to a specific methodological layer.

5.3 Robustness and ablation tests

The predictive controller should also be tested under forecast degradation, alternative battery limits, hydro availability reductions, and shifted evening peaks. A particularly informative ablation is to compare perfect-forecast scheduling with realistic forecast-driven scheduling, because it reveals how much of the apparent optimization benefit is dependent on information that an operator would not actually have. This is consistent with recent work showing that forecast uncertainty can materially change hydro and battery operating behavior (Phillips et al., 2026).

6. Results and Engineering Interpretation

The present study's empirical evidence is limited to the documented baseline simulation, not a completed AI experiment. The baseline nevertheless provides a meaningful engineering result: 82.98% renewable supply is achieved with a simple dispatch hierarchy, but 180.57 MWh of the 1,061.00 MWh daily demand still relies on fossil fallback. The battery finishes at 7.00 MWh, equivalent to 10% of the stated 70 MWh system capacity, after delivering 53.28 MWh (Arpit, 2026).

The pattern suggests a temporal mismatch rather than an absolute lack of renewable assets. The system has significant hydro contribution and some solar-to-battery shifting, yet the final state of charge is low and fossil fallback remains material. In the baseline architecture, the controller reacts to deficits after they occur. A predictive controller would instead reserve part of the battery and hydro flexibility against forecasted peaks. Whether that strategy actually lowers fossil use must be demonstrated experimentally; it cannot be inferred from the baseline alone.

This distinction is also important for publication quality. The 82.98% renewable share should be reported as a baseline simulation result, while improvements from AI, optimization, or forecasting should be reported only after an independent experiment produces them. Keeping those claims separate strengthens the scientific contribution because the reader can identify exactly which statements are measured and which are proposed.

7. Smart Sustainable Community Integration

The smart-community significance of the framework lies in the conversion of local energy assets into an interoperable, data-driven service layer. ITU-T Recommendation Y.4614, approved in 2026, defines an energy data model for a city-level energy management platform and specifies requirements for energy data sharing to support integrated city energy services. ITU-T Recommendation Y.4242, also approved in 2026, addresses requirements and capabilities for energy-storage services in residential smart-city communities (International Telecommunication Union, 2026a, 2026b). Together, these recommendations create a useful standards context for a predictive EMS: forecasting variables, battery states, dispatch decisions, and performance indicators should be represented as structured, shareable energy information rather than isolated outputs from a single script.

The proposed framework also fits the logic of ITU-T Recommendation Y.4903, which provides key performance indicators and selection principles for assessing smart sustainable cities across economic, environmental, and social dimensions (International Telecommunication Union, 2022). For this study, the most immediate indicators are environmental and infrastructure-oriented: renewable supply share, fossil dependence, storage utilization, carbon footprint, and service reliability. A future city deployment could extend the same framework to cost, energy affordability, outage resilience, and equitable access.

The standards mapping is architectural rather than a claim of formal compliance or certification. Accordingly, Table 7 maps the proposed EMS variables, outputs, and KPIs to the relevant smart-community energy-management functions.

7.1 Proposed AI-energy KPI set for interoperable evaluation

- Forecast quality: mean absolute error (MAE), root mean square error (RMSE), and normalized error for load and PV forecasts.
- Dispatch quality: renewable supply share, fossil energy, PV curtailment, and power-balance residual.
- Storage quality: SOC compliance, equivalent full cycles, terminal reserve, and degradation proxy.
- Hydro utilization: renewable balancing energy, reserve availability, and maximum operating-limit violations.
- Sustainability: CO_2 emissions associated with residual fossil supply.
- Community resilience: unmet demand duration, recovery time, and reserve adequacy under disturbed conditions.

Table 7.

Mapping of proposed EMS data fields and KPIs to smart-community energy services.

EMS variable/output	Example fields	Smart-community function	KPI

EMS variable/output	Example fields	Smart-community function	KPI
Load/PV forecast	Predicted power, horizon, forecast error	Anticipatory energy management	MAE/RMSE
BESS state	SOC, charge/discharge power, availability	Energy-storage service	SOC compliance, throughput
Hydro state	Available reserve, dispatch power, availability	Renewable balancing	Reserve adequacy
EMS action	PV, BESS, hydro, fallback commands	Coordinated dispatch	Power-balance residual
Environmental output	Renewable share, fossil energy, CO_2	Sustainability reporting	Carbon reduction
Reliability output	Unmet demand, duration, reserve margin	Community resilience	Reliability index

8. Discussion

The most important design decision is to treat AI as a forecasting and decision-support component rather than as an unrestricted controller. In energy systems, an accurate prediction does not guarantee a safe dispatch, and a powerful optimizer can still fail when its assumptions do not match operating reality. The layered design therefore separates estimation, decision making, and feasibility enforcement.

The second design principle is reversibility. The original rule-based controller remains available as a fallback mode. This enables a practical migration path from an educational prototype to an experimental predictive EMS: first validate forecasts, then validate constrained optimization, then test bounded learning correction. Such a progression reduces the risk of attributing improvements to AI when they actually come from better data, different weights, or relaxed physical assumptions.

The third principle is reproducibility. The GitHub repository exposes the conceptual model, stated parameters, result metrics, output structure, and limitations (Arpit, 2026). A journal version should preserve this openness by publishing the exact dataset or assumptions used for each experiment, fixed train-test splits, model hyperparameters, random seeds, constraint definitions, and evaluation scripts. Without those details, an impressive renewable-share number is difficult to reproduce and compare.

8.1 Implications for deployment

A production deployment would require substantially more information than the current prototype contains, including validated plant operating curves, time-stamped load and PV measurements, battery efficiency and degradation characteristics, hydraulic constraints, reserve rules, network limits, and a defined market or tariff context. The present manuscript therefore supports a decision-support research prototype, not direct control of the Rani Avanti Bai plant.

Before this framework is presented as an experimental AI paper, the minimum evidence package should include: a documented multi-day or seasonal load/PV dataset; at least one independently evaluated forecasting model; a rolling-horizon constrained optimizer using those forecasts; an apples-to-apples comparison against the existing rule-based controller; sensitivity tests for SOC, hydro availability, forecast error, and peak demand; and complete reporting of model parameters, objective weights, solver settings, and random seeds. Only after this package is executed should conclusions such as reduced fossil fallback, improved renewable utilization, or improved reliability be stated as experimental findings.

9. Limitations and Future Research

The study has five material limitations. First, the current evidence is a 24-hour simulation rather than a seasonal or annual assessment. Second, the PV and demand profiles are synthetic rather than measured plant datasets. Third, the hydropower block is simplified and does not reproduce hydraulic dynamics, unit commitment, or water-value optimization. Fourth, no trained AI forecasting experiment is included in the present baseline. Fifth, although the baseline parameters are now explicitly documented, the emission-factor source and empirical justification for the selected battery and synthetic profile assumptions should be archived with the journal release.

These limitations define the next research sequence. The immediate step is to import a reproducible time series for load and solar generation, train and compare candidate forecasting models, and evaluate them with a rolling-origin protocol. The second step is to implement the constrained rolling-horizon optimizer on the same time series. The third step is to extend the study to monthly and annual scenarios, with wet, normal, and dry hydrological availability, and add battery degradation and economic objectives. Only after those stages should a learning-based correction policy be compared with the optimization baseline.

10. Conclusion

This paper has reformulated the Jabalpur Hydro–PV–BESS simulation from a static source-priority exercise into a research framework for predictive energy management. The documented baseline already demonstrates a substantial renewable contribution of 82.98% of the simulated daily supply, but still leaves 180.57 MWh to fossil fallback and ends with a low battery state of charge (Arpit, 2026). The proposed Forecast–Coordinate–Constrain architecture addresses the underlying control problem by combining short-horizon prediction, coordinated use of PV, BESS, and hydropower, and explicit physical constraints. Its contribution is intentionally methodological: it provides a transparent route for moving from an interpretable rule-based prototype to an empirically testable AI-enabled EMS. The framework also aligns naturally with the emerging smart-city energy data and storage-service standards represented by ITU-T Recommendations Y.4614 and Y.4242, while using Y.4903 as a guide for measurable sustainability-oriented performance indicators (International Telecommunication Union, 2026a, 2026b, 2022). The next stage of work should therefore focus on reproducible time-series data, forecast validation, rolling-horizon optimization, degradation-aware storage operation, and multi-season testing. That progression would allow AI-related gains to be reported as measured engineering evidence rather than assumed performance, strengthening the basis for peer-reviewed publication.

11. Data and Code Availability

The underlying prototype and its documented assumptions are publicly maintained in Oihika Arpit's GitHub repository. The repository README describes the work as an offline educational simulation and explicitly states that it does not connect to real-time

SCADA, PLC, turbine, gate, inverter, breaker, battery controller, or electrical equipment (Arpit, 2026). The baseline parameter values used for the reported experiment are documented in the manuscript and in the supplied Python configuration. The journal release should preserve the exact code, input profiles, configuration, and emission-factor source alongside the reported results.

For direct access to the executable baseline and supporting files, the repository is also provided as a QR code in the original manuscript. The QR code is a navigation aid; reproducibility depends on the documented code, input profiles, configuration values, and generated result files.

Figure 1. QR code linking to the public GitHub repository for the Hydro–PV–BESS baseline simulation.

12. Author Contributions

Oihika Arpit: Conceptualization; methodology; hybrid energy-system modelling; simulation design; analysis and interpretation; predictive-framework design; visualization; manuscript drafting.

Jasreet Kaur: Literature review and synthesis; manuscript review and editing; contribution to research positioning and scholarly presentation.

Dr. Smrite Goudhaman: Supervision; research guidance; methodological review; critical review and editing; oversight of the manuscript's academic positioning.

13. AI-Assisted Manuscript Preparation Disclosure

Generative AI tools were used during manuscript preparation for language editing, structural refinement, and drafting support. The authors retained responsibility for the technical concept, project parameters, interpretation of simulation results, equations, methodological decisions, source selection, and final manuscript content. No AI-generated numerical result is presented as an experimentally measured result unless it is traceable to the documented project simulation or an independently verified source.

14. References

Arpit, O. (2026). *Hydro–PV–BESS energy management system: Jabalpur hydro–solar–battery hybrid energy simulation* [Computer software]. GitHub. https://github.com/oihika/Hydro-PV-BESS-Energy-Management-System-JABALPUR-RANI-AVANTI-BAI-HYDEL-POWER-PLANT-2-45-

Casazola-Cruz, O., Pilco-Nuñez, A., & Arroyo-Paz, A. (2025). Machine learning for photovoltaic power forecasting integrated with energy storage systems: A scientometric analysis, systematic review, and meta-analysis. *Energies, 18*(23), 6291. https://doi.org/10.3390/en18236291

Central Electricity Authority. (2024). *Final report on resource adequacy plan for the state of Madhya Pradesh (2024–25 to 2034–35).* Ministry of Power, Government of India.

Grimaldi, A., Minuto, F. D., Brouwer, J., & Lanzini, A. (2024). Optimal energy management of a utility-scale battery energy storage system integrated with a photovoltaic power plant considering battery degradation. In *2024 IEEE Power & Energy Society General Meeting (PESGM)*. https://doi.org/10.1109/PESGM51994.2024.10761082

Gupta, M., Arya, A., Varshney, U., Mittal, J., & Tomar, A. (2025). A review of PV power forecasting using machine learning techniques. *Power & Energy Systems*, 100058. https://doi.org/10.1016/j.pes.2025.100058

International Telecommunication Union. (2022). *ITU-T Recommendation Y.4903 (03/2022): Key performance indicators for smart sustainable cities to assess the achievement of sustainable development goals.*

International Telecommunication Union. (2026a). *ITU-T Recommendation Y.4242 (05/2026): Requirements and capability framework of an energy storage service for residential community in the smart city.*

International Telecommunication Union. (2026b). *ITU-T Recommendation Y.4614 (05/2026): Energy data model for a city-level energy management platform.*

Madhya Pradesh Power Generating Company Limited. (2026). *Comprehensive renovation modernization and uprating of 2×45 MW Rani Avanti Bai Sagar Hydro Power Station, Bargi, District Jabalpur, Madhya Pradesh* (Government of Madhya Pradesh e-Procurement Tender 2025_MPPGC_461283_1).

Pengtem, C., Deeum, S., Amirullah, Ohgaki, H., Romphochai, S., Bhumkittipich, P., & Bhumkittipich, K. (2025). Design of multi-objective energy management for remote communities connected with an optimal hybrid integrated photovoltaic–hydropower–battery energy storage system (PV-HP-BESS) using improved particle swarm optimization. *Energies, 18*(9), 2250. https://doi.org/10.3390/en18092250

Phillips, T., Lopez, C. J., Kwon, J., Mahalik, M., Durvasulu, V., & Balliet, W. H. (2026). Improving the realism of hydropower and hydropower–battery hybrid plant optimization: A forecast-aware rolling horizon model. *Renewable Energy, 273*, 126109. https://doi.org/10.1016/j.renene.2026.126109

Tang, W., Huang, Y., Qian, T., Wei, C., & Wu, J. (2025). Coordinated central-local control strategy for voltage management in PV-integrated distribution networks considering energy storage degradation. *Applied Energy, 389*, 125684. https://doi.org/10.1016/j.apenergy.2025.125684

Zambrano, H. H. N., Benigni, A., Lazzari, R., & Piegari, L. (2025). A forecast-driven energy management methodology for microgrids under seasonal variability. In *2025 International Conference on Clean Electrical Power (ICCEP).*

Zhu, Y., et al. (2025). A review on energy management system for grid-connected utility-scale renewable hybrid power plants. *WIREs Energy and Environment*, e70004. https://doi.org/10.1002/wene.70004