FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 10 No. 21, November, 2026, pp 303 – 317
DOI: https://doi.org/10.33003/fjs-2026-1021-6089

AI-SEHOS: A Solar-Powered, Energy-Aware Digital Learning Framework for Nomadic Communities in Northern Nigeria

[1]Gregory Onwodi, [1]Enoch Jacob Dodo, [1]Grace Egenti and [2]Yayock Amos Takai

[1]Department of Information System and Technology, Faculty of Computing, National Open University of Nigeria, Jabi, FCT Abuja, Nigeria.
[2]Department of Information Technology, Faculty of Computing, Federal University of Applied Sciences Kachia, Kaduna, Nigeria.

*Corresponding authors' email: enochdodo46@gmail.com Phone: +234708 2241514

Received: 14 Aug. 2026
Accepted: 19 Oct. 2026
Published: 05 Oct. 2026

ABSTRACT
Nomadic and pastoralist communities in Northern Nigeria are poorly served by digital-education strategies because electricity is absent, households are seasonally mobile and ICT infrastructure is weak. This concept paper specifies AI-SEHOS (AI-driven Solar Energy Harvesting, Optimisation and Storage), a four-layer framework (Data, Energy, AI Decision and Educational Service) that treats instructional demand and mobility as explicit terms in energy dispatch and load allocation. The specification uses twenty-five numbered equations adapted from established photovoltaic, forecasting, battery-estimation and reinforcement-learning work. The components were exercised on a documented synthetic generator with chronological splits, repeated runs and stated baselines; the code is supplied. On synthetic data only, long short-term memory (LSTM) and gradient-boosting forecasters reduced one-hour-ahead test RMSE by about 4–6% relative to smart persistence (nRMSE 2.52% and 2.56% of capacity versus 2.67%) and were statistically indistinguishable; the gain disappeared under a simpler generator. A tabular Q-learning dispatch policy matched, but did not beat, the best rule-based controller on 30 held-out traces. Criticality-aware allocation raised server availability during school hours from 81% to 98% and reduced weighted unmet load by about 34% under fixed reference weights, at the cost of more lighting shortfall; the allocation depended only on weight rank. These results illustrate mechanisms; they are not evidence of field performance. The paper contributes a specification, a reproducible test bed, a field-validation protocol with pre-stated hypotheses and an explicit account of what remains unestablished.

Keywords: Solar Energy Management, Battery Energy Storage, Reinforcement Learning, LSTM Forecasting, Nomadic Education, Off-Grid Digital Learning, Northern Nigeria

INTRODUCTION

The global shift toward renewable energy reflects rising energy demand, climate pressure and the limits of fossil-fuel supply. Solar, wind, biomass, hydropower and geothermal technologies now underpin decentralised electrification where extending the grid is technically or economically infeasible (Ukoba et al., 2024). Education is among the sectors most affected by this deficit. Moner-Girona et al. (2025), who mapped schools across the African continent, estimate that 32% of African school-aged children live nearest to a school without electricity, a condition that constrains digital learning and instructional quality.

In Northern Nigeria the deficit is compounded for nomadic and pastoralist populations, predominantly Fulani communities, whose seasonal migration has historically placed them beyond the reach of fixed-site schooling (Aderinoye et al., 2007). Nomadic-education interventions coordinated by the National Commission for Nomadic Education, including mobile schools and radio-based instruction, have widened access but remain limited by weak ICT and power infrastructure (Aderinoye et al., 2007; Olaniran, 2018). Where generators are used, fuel is a recurrent cost and a logistical burden for mobile sites. Station-based estimates place the annual mean global horizontal irradiation of the northern states at roughly 5–6 kWh m^{-2} day^{-1}, with Sokoto at about 5.9 kWh m^{-2} day^{-1} (Bawonda et al., 2023), which makes photovoltaic (PV) generation the most plausible energy source for mobile learning infrastructure. The resource is, however, seasonal: harmattan haze, rainy-season cloud, high module temperatures and relocation of equipment all reduce and interrupt yield. Mobile pastoralist households are already adopters of small-scale solar technology, with solar panels on nomadic tents powering audiovisual equipment, which suggests that the barrier is the absence of managed, education-oriented systems rather than of acceptance of solar power as such (Krätli & Dyer, 2009).

Artificial intelligence (AI) has meanwhile become a core capability in renewable-energy engineering. Machine-learning (ML), deep-learning (DL) and reinforcement-learning (RL) methods are used for irradiance and demand forecasting, maximum-power-point tracking (MPPT), battery management and load scheduling (Ejiyi et al., 2025; Khouili et al., 2025; Ukoba et al., 2024), building on well-established device-level PV models (Song et al., 2021) and MPPT practice (Worku et al., 2023). A largely separate literature documents AI in education, including adaptive pathways, tutoring, assessment and learning analytics (Alfredo et al., 2024). Within the sources examined for this paper (Section 2.7 states the limits of that review), we found no treatment that links the two: a quantitatively specified system that forecasts and manages solar generation and storage while allocating the resulting energy budget against the instructional needs of a mobile, off-grid learner population.

This paper therefore proposes and specifies AI-SEHOS, a four-layer, learner-aware architecture intended to sustain digital learning centres for nomadic communities. It is a concept paper: its contribution is the specification, its formalisation, a transparent synthetic test of its mechanisms and a protocol for field validation. It is not an empirical test of field performance. The paper addresses four research questions.

Q1. What do the examined literatures establish about AI-based solar forecasting, battery-state estimation and storage dispatch, and what do they leave unaddressed for mobile, off-grid education? Q2. How can instructional demand, mobility and seasonal solar conditions be represented explicitly in an energy-management specification? Q3. Under a documented synthetic scenario with held-out evaluation, how do the forecasting, dispatch and load-allocation components behave relative to simple baselines, and which claims do the results support? Q4. What evaluation protocol, with what pre-stated hypotheses, would allow field validation of the framework?

Literature Review

This section reviews seven strands of work: AI in renewable energy; PV device modelling and MPPT; solar forecasting and battery management; RL-based microgrid dispatch; rural electrification and school access; AI in education; and nomadic education in Nigeria. The review is purposive rather than systematic: sources were chosen because they supply the models, benchmarks or context on which the framework relies, and Section 2.7 sets out what that approach can and cannot support.

Artificial Intelligence in Renewable Energy Systems

Ejiyi et al. (2025) reviewed the literature on AI in renewable energy and identified ML, DL and RL as the dominant technique families, with reported gains in forecasting, PV optimisation, storage management, predictive maintenance and smart-grid stabilisation; they also highlight the challenges that hinder adoption in practice. Ukoba et al. (2024) similarly report that ML and neural-network methods improve resource assessment, forecasting, monitoring and grid-integration control, while cautioning that data variability, limited interpretability and difficulty of real-time adaptation remain open problems, particularly outside data-rich, grid-connected contexts. These reviews support the technical premise of this paper, that AI can improve renewable-energy performance in the settings studied, while leaving open whether the gains transfer to small, mobile, education-serving microsystems.

PV Modelling, Solar Forecasting and Battery Energy Storage

Device-level representation is a precondition for simulation-based design. Song et al. (2021) present a non-iterative extraction method for the five-parameter single-diode model, validated against datasheets across several PV technologies. Module output also depends on temperature, with power falling as cell temperature rises (Skoplaki & Palyvos, 2009), which matters in a region where module temperatures are high. Worku et al. (2023) review conventional (perturb-and-observe, incremental conductance), soft-computing and hybrid MPPT techniques under partial shading and conclude that hybrid strategies track quickly and efficiently but are complex to implement, which is relevant to low-cost controllers in mobile kits.

Short-horizon forecasting is a second precondition for intelligent dispatch. Solar-forecasting reviews establish clear-sky-referenced (smart) persistence as the standard naive benchmark, because raw persistence is penalised by the diurnal cycle (Inman et al., 2013; Yang et al., 2018). Khouili et al. (2025) systematically reviewed DL-based PV forecasting and report that ambient temperature, pressure and humidity are the most common input features, and that robustness and interpretability remain priorities. The LSTM recurrence itself descends from Hochreiter and Schmidhuber (1997), with the forget gate used in modern formulations introduced by Gers et al. (2000). Gradient-boosted trees (Chen & Guestrin, 2016) are a strong tabular alternative and are included here as a baseline.

On the storage side, state-of-charge (SOC) estimation underpins safe dispatch. Open-loop Coulomb counting drifts under sensor noise and unknown initial conditions, so filter-based correction is standard: the Kalman filter (Kalman, 1960), its extended form applied to lithium batteries (Plett, 2004) and the unscented filter (Julier & Uhlmann, 2004), which avoids linearising the nonlinear open-circuit-voltage relation. Hu et al. (2022) evaluate Kalman-filter algorithms for lithium-battery SOC estimation and report improved accuracy over simpler estimators under the conditions they test. Battery ageing is a further concern in hot climates: Leng et al. (2015) show that the cyclic ageing rate of lithium-ion cells increases when they operate above room temperature, in the range of 25–55 °C they examined, which motivates the temperature-aware charge limits and ventilated housings proposed in Table 2.

For dispatch, Alabdullah and Abido (2022) formulate microgrid energy management as a Markov decision process solved with a deep Q-network, and Mohammadi et al. (2024) review control strategies and report growing use of RL because it avoids repeated re-optimisation as conditions change. Deep Q-networks (Mnih et al., 2015) and deterministic policy-gradient methods such as DDPG (Lillicrap et al., 2016) are the usual function-approximation routes, and tabular Q-learning (Watkins & Dayan, 1992) is the baseline from which they derive. Demand-side management, in which loads are shifted or shed according to supply, is the conceptual precursor of criticality-aware allocation (Palensky & Dietrich, 2011). Most of the RL evidence comes from electric-mobility and grid-scale contexts, not off-grid educational microgrids, so its transfer is an assumption that must be tested.

Finally, running forecasting and control on the device itself is an embedded-computing problem. Reviews of tiny machine learning describe how compact models can run on microcontrollers within severe memory and power budgets (Alajlan & Ibrahim, 2022; Ray, 2022), which is the relevant regime for an offline, solar-powered learning centre and is taken up in Section 7.4.

Rural Electrification, Solar Mini-Grids and School Access

Carabajal et al. (2024) report a cohort study of households and businesses connected to solar mini-grids in Kenya and Nigeria and find significant increases in economic activity and productivity among connected households and businesses. The study did not examine AI-based energy management or digital-learning integration. Moner-Girona et al. (2025), in a continent-wide geospatial analysis of African schools, estimate that 32% of African school-aged children live nearest to an unelectrified school and that decentralised solar PV could electrify the schools that lack power at an upfront cost of about €2 billion. Their analysis concerns fixed-site schools, not mobile learning centres whose demand shifts geographically and seasonally with migration. Moner-Girona et al. also note that evaluations of school electrification, including Koima (2024) in rural Kenya, have almost uniformly reported positive academic outcomes, and Kebir et

al. (2023) examined second-life battery systems as an affordable route to energy access in Kenyan primary schools. Together these studies establish the case for solar power in schools; none addresses managed, learner-aware energy allocation for mobile sites.

Artificial Intelligence in Digital Education

Alfredo et al. (2024), reviewing empirical studies of human-centred learning analytics and AI in education, report that AI-supported analytics improves early identification of at-risk learners and enables adaptive content, while noting limited learner and teacher participation in system design and under-examined transparency and oversight. The reviewed systems were largely developed in institutions with stable electricity and connectivity, an assumption that does not hold in nomadic settings. Technology adoption in such settings should also be examined through established adoption frameworks, namely the Technology Acceptance Model (Davis, 1989) and Diffusion of Innovations (Rogers, 2003), which Section 6 uses to specify adoption indicators.

Nomadic Education in Northern Nigeria

Nigeria's Nomadic Education Programme has for decades sought to extend basic education to pastoral Fulani, migrant fishing and other mobile communities through mobile schools, radio-based instruction and itinerant teaching (Aderinoye et al., 2007). Aderinoye et al. (2007) document inadequate infrastructure, teacher shortages tied to staffing mobile schools and a mismatch between rigid curricula and pastoral livelihoods, and identify mobile and distance technologies as a promising but under-developed route. Krätli and Dyer (2009) argue more broadly that education for mobile pastoralists should be fitted to mobile livelihoods rather than imposed through fixed-site schooling. Olaniran (2018) describes how the National Commission for Nomadic Education has struggled to turn policy into practice for the related Almajiri population, citing underfunding, weak inter-agency coordination and curriculum misalignment, and Umar and Maaji (2025) show that gaps in nomadic and Almajiri education continue to compound public-health, child-protection and security risks. These sources establish that access barriers persist; they do not themselves measure present-day power availability at nomadic learning sites, and current programme-level data from Nigerian agencies and international organisations should be obtained before field work (Section 3.2).

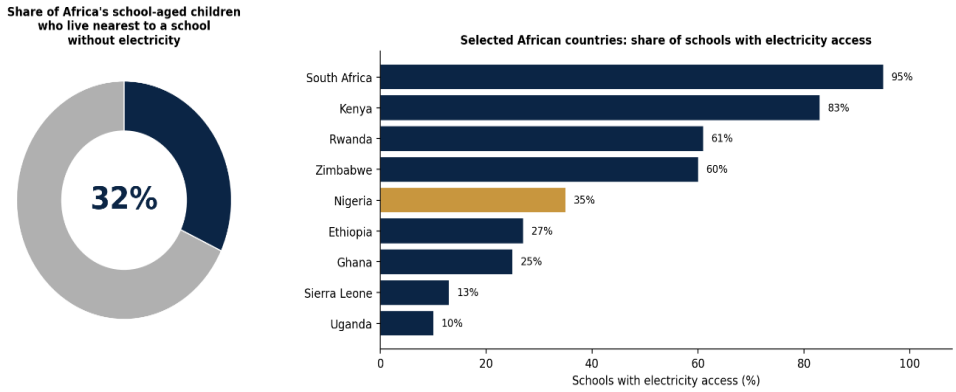


Figure 1: Electricity Access in African Schools: the Share of School-Aged Children who live Nearest to a School Without Electricity (left) and the Share of Schools with Electricity Access in Selected Countries, with Nigeria Highlighted (Right). Redrawn by the Authors from Moner-Girona et al. (2025; Table 1 and Summary). Country Values are the Shares of Schools with electricity Reported in Table 1 of the Source, which Draws on National Statistics and Giga Datasets where Available and on a Grid and Night-Light Proxy Elsewhere

Synthesis and Research Gap

Table 1 summarises the seven strands. Each supports one component of the problem. Within the sources examined, none combines (i) forecasting and optimisation of intermittent solar generation at device and time-series level, (ii) battery management with formal state estimation under classroom-linked, mobility-affected demand and (iii) allocation of the resulting energy budget, through an explicit objective function, to sustain digital instruction for a geographically mobile learner population. The novelty claimed here is therefore architectural and limited: the individual equations in Section 4 are adapted from established formulations, and the contribution is their integration, specifically the treatment of instructional demand and mobility as explicit terms in the dispatch and load-allocation objectives (Equations 14, 15 and 18), together with a transparent synthetic test bed and a field-validation protocol.

Table 1: Synthesis of the Reviewed Literature and the Corresponding Contribution of this Study

Strand	Representative source(s)	What the source supports	Gap addressed by AI-SEHOS
AI and renewable energy	Ejiyi et al. (2025); Ukoba et al. (2024)	ML/DL/RL improve forecasting, optimisation and storage management, mainly in industrial and utility-scale settings.	Extends AI energy management to a decentralised, education-serving microgrid.
PV device modelling and MPPT	Song et al. (2021); Worku et al. (2023);	Single-diode extraction, temperature dependence and hybrid	Device- and controller-level equations (Eqs. 2–5) with

Strand	Representative source(s)	What the source supports	Gap addressed by AI-SEHOS
	Skoplaki & Palyvos (2009)	MPPT improve modelled and harvested energy.	temperature derating and relocation downtime.
Solar forecasting and battery management	Khouili et al. (2025); Hu et al. (2022); Inman et al. (2013); Leng et al. (2015)	LSTM/CNN models are widely used for short-horizon PV forecasting; Kalman-type filters correct SOC estimates; heat accelerates cell ageing.	LSTM forecasting and filter-corrected SOC (Eqs. 6–13) inside a learner-demand-aware loop.
RL for microgrid dispatch and demand-side management	Alabdullah & Abido (2022); Mohammadi et al. (2024); Palensky & Dietrich (2011)	DQN/MDP dispatch is reported to outperform rule-based and classical control in microgrid studies; loads can be shifted or shed according to supply.	MDP dispatch with instructional demand and a relocation flag (Eqs. 14–17); tested against rule baselines.
Rural electrification and school access	Carabajal et al. (2024); Moner-Girona et al. (2025); Koima (2024); Kebir et al. (2023)	Solar mini-grids and school electrification improve economic and educational outcomes at household, national and continental scale.	Targets mobile, seasonally migrating learning centres, not fixed sites.
AI in education	Alfredo et al. (2024)	Learning analytics improve personalisation and at-risk identification, assuming stable power and ICT.	Couples the educational service to an explicitly managed energy budget.
Nomadic education, Nigeria	Aderinoye et al. (2007); Krätli & Dyer (2009); Olaniran (2018); Umar & Maaji (2025)	Infrastructure, staffing and curriculum barriers persist; mobile technologies are promising.	Proposes an AI-managed energy layer for mobile centres; current site-level evidence still needed.

Scope and Limits of this Review

The review is purposive, not systematic. No database search strings, date limits or screening procedure were used, and the base of sources is small and weighted toward recent review articles. Statements such as "within the sources examined, no treatment links these strands" are therefore bounded by that base and are not claims that no such work exists, and an absence within a purposive selection is not evidence of global novelty. Relative to an earlier draft, the review now also draws on school-electrification, embedded-computing, battery-ageing, demand-side-management and pastoralist-education sources, but it remains a sample of these literatures: solar-powered ICT projects and offline-first educational technology in Africa, and recent programme-level evidence on nomadic education in Nigeria, are still only partly covered. A registered systematic mapping covering these areas is listed as future work (Section 8) and should precede any stronger novelty claim.

MATERIALS AND METHODS

Methodology

Research Philosophy and Design Science Research Approach

Because the aim is to design and evaluate an artefact, an AI-managed renewable-energy and educational-service framework, rather than to explain an existing phenomenon, the study adopts a pragmatist stance and Design Science Research (DSR) (Hevner et al., 2004; Peffers et al., 2007). Pragmatism accommodates the combination of quantitative evidence (irradiance, PV output, battery telemetry, forecasting accuracy) with qualitative stakeholder evidence (teachers, caregivers, nomadic-education officers, energy engineers) that the evaluation needs. Following Peffers et al. (2007), DSR comprises six activities, shown in Figure 2: problem identification and motivation, definition of objectives, design and development, demonstration, evaluation and communication. This paper completes the first two, completes the specification part of design and development together with synthetic-data tests of its mechanisms, and communicates the result. Demonstration and evaluation in the field are specified as a protocol in Section 6 and have not been executed.

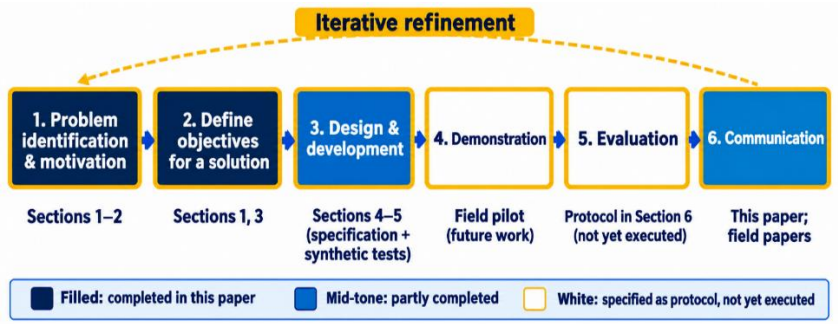


Figure 2: Design Science Research Activities (Peffers et al., 2007) and their Status in this Paper

Study Context, Prospective Field Protocol and Ethics

The framework is contextualised for the Sudano-Sahelian states of Northern Nigeria, such as Kano, Katsina, Sokoto, Borno, Yobe and Bauchi, because of their solar resource and nomadic and pastoral populations. State selection for any pilot should be justified with current data on nomadic populations and mobility corridors. Many pastoralist routes also extend into the Middle Belt, and security conditions in several northern states may restrict field access and equipment safety. These factors will determine pilot sites, and Table 2 lists them as design inputs.

The field work described in this subsection and in Section 6 is prospective: it is specified, not performed. The planned mixed-methods pilot uses purposive sampling of six to eight nomadic learning centres across at least two ecological zones, with semi-structured interviews with teachers, caregivers and programme officers (target 12–20 interviews, extended until thematic saturation is judged to be reached), conducted in Hausa or Fulfulde with interpreters where needed. The interview target is justified by empirical studies of saturation, which found that saturation was typically reached within 9–17 interviews or 4–8 focus-group discussions in relatively homogeneous samples (Hennink & Kaiser, 2022). The six-to-eight-site design is a feasibility-scale pilot, not a powered confirmatory trial: its purpose is to test procedures, estimate variability and show whether the energy, service and adoption indicators can be measured, in line with the recommendation that pilot studies be designed to assess feasibility rather than to test hypotheses (Eldridge et al., 2016; Lancaster et al., 2004). The hypotheses in Table 8 are accordingly estimation targets with prospective thresholds, and any confirmatory trial would require a sample size calculated from the pilot variance estimates.

The comparison strategy follows from the finding of Section 5.3 that a learned controller must be shown to outperform simple rules. At each site the AI-managed controller and the best rule-based controller would operate in alternating blocks of seven days, with block order randomised per site, so that each site serves as its own control and weather, load and relocation patterns are balanced across arms; sites would be analysed with paired estimates and confidence intervals. Quantitative inputs would be site irradiance and temperature, battery telemetry, attendance and device-usage logs, and secondary resource data (NASA, n.d.; World Bank Group, n.d.; Nigerian Meteorological Agency records). No data collection will begin before approval by the university research ethics committee and community and gatekeeper consent, with caregiver consent and child assent for any learner-linked data. Attendance and device-usage telemetry constitute personal data and will be handled under the Nigeria Data Protection Act, 2023, using data minimisation, local-first storage and pseudonymisation. Because this paper reports only synthetic simulations, no human-subject data were collected and no ethics approval was required for the work reported.

Deployment Constraints Treated as Design Inputs

Mobility is the defining feature of the target population, so it is represented explicitly. In the model, a relocation day (a) removes classroom load for that day and (b) removes PV generation during transport and set-up hours, and a binary relocation flag m_t enters the dispatch state (Section 4.3.4). Seasonality enters through weather regimes (clear, harmattan haze, partly cloudy, overcast) whose frequencies vary by season, and temperature enters through cell-temperature derating (Equation 4). Other constraints that cannot yet be quantified are listed in Table 2 with the design response adopted.

Table 2: Deployment Constraints and the Corresponding Design Responses (to be Validated in the Field)

Constraint	Effect on the system	Design response in AI-SEHOS
Relocation and seasonal migration	Loss of generation and class days during transport; changes in panel orientation and shading.	Relocation flag in state; zero-PV hours in simulation; ruggedised, quick-deploy mounting; re-calibration of forecaster after relocation.
Harmattan haze and dust	Lower clear-sky index; soiling losses.	Haze regime in generator; scheduled cleaning; forecaster trained on seasonal data.
High ambient and module temperature	PV derating; battery ageing and thermal-safety risk (Leng et al., 2015).	Temperature term in Equation 4; temperature-aware charge limits; shaded, ventilated battery housing.
Vibration and transport	Mechanical damage to panels, connectors and cells.	Protective cases, secured cabling, battery management system with fault logging.
Theft and security	Loss of equipment; restricted access to some areas.	Community custody agreements, non-resale-friendly components, site selection reflecting security advice.
Intermittent connectivity	Cloud analytics unavailable.	Offline-first learning content; local inference; periodic synchronisation.
Limited local maintenance capacity	Prolonged downtime after faults.	Training of community custodians; simple fault indicators; remote diagnostics when connected.

The AI-SEHOS Framework

System Architecture

AI-SEHOS has four layers coordinated through a feedback loop (Figure 3). The Data Layer aggregates environmental sensors (irradiance, ambient and panel temperature, humidity), an energy meter, battery telemetry (voltage, current, SOC, state of health, temperature) and educational metadata (attendance, device use, schedule, relocation plans) through a local gateway with local and, when available, cloud storage. The Energy Layer comprises the PV array, charge controller, lithium battery bank, inverter and loads. The AI Decision Layer performs forecasting, battery-state estimation, dispatch and load allocation. The Educational Service Layer hosts an offline-first learning-management system and content, learner devices and, connectivity permitting, cloud-linked analytics of the kind reviewed by Alfredo et al. (2024). Instruction is thereby decoupled from connectivity but, by design, coupled to the availability of solar energy.

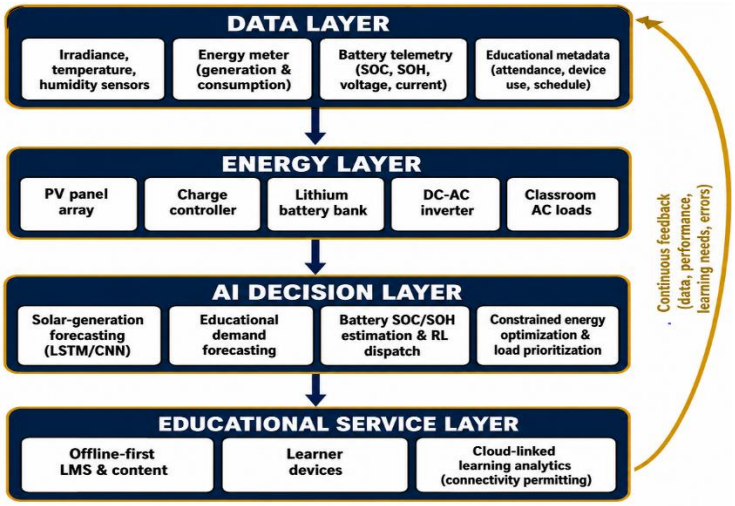


Figure 3: AI-SEHOS Four-Layer Architecture with Continuous Feedback

Table 3 defines the notation used in the remainder of this section, with units, and states which components are simulated in Section 5 and which are specified only. Power is expressed in W, energy in Wh, and the time step is $\Delta t = 1$ h throughout, so that energy equals power multiplied by Δt.

Table 3: Notation, Units and Implementation Status, Organised by Governing Equation

Symbols	Definition (units)	Status in this paper
E, P, PSH, η	Daily energy yield (Wh day^{-1}), rated array power (W), peak sun hours (h day^{-1}), system efficiency (–) (Eq. 1)	Sizing relation; not simulated
I, V, I_{ph}, I_0, n, N_s, R_s, R_{sh}, V_T, k, T, q	Cell or module current (A) and voltage (V), photocurrent (A), saturation current (A), ideality factor (–), cells in series (–), series and shunt resistance (Ω), thermal voltage (V), Boltzmann constant (J K^{-1}), absolute temperature (K), electron charge (C) (Eqs. 2–3)	Specified; not simulated
P_{pv}, P_{stc}, G, G_{stc}, γ, T_c, T_a, k_G, T_{ref}, η_{sys}	PV power (W), power at standard test conditions (W), irradiance (W m^{-2}), reference irradiance (1000 W m^{-2}), temperature coefficient (°C^{-1}), cell and ambient temperature (°C), irradiance-to-temperature coefficient (°C m^2 W^{-1}), reference temperature (25 °C), lumped system efficiency (–) (Eq. 4)	Simulated (generator)
dP/dV	Slope of the power–voltage curve (W V^{-1}); zero at the maximum power point (Eq. 5)	Specified; not simulated (effect subsumed in η_{sys})
x_t, h_{t-1}, f_t, i_t, $\tilde{c}_t$, c_t, o_t, h_t, σ, $\odot$, $W_\cdot$, $b_\cdot$	LSTM input vector, previous hidden state, forget, input and output gates, candidate and cell states, hidden state, sigmoid, Hadamard product, weight matrices and bias vectors (all –) (Eqs. 6–10)	Simulated
$\hat{E}_{t+1}$, g, θ	One-step-ahead generation forecast (W), read-out function and its parameters (Eq. 11)	Simulated
SOC, C_n, η_c, i, τ	State of charge (–), nominal capacity (Ah), coulombic efficiency (–; below one when charging, one when discharging), current (A, positive when charging), integration time (s) (Eq. 12)	Simulated as noise-free energy bookkeeping
x_k, u_k, z_k, w_k, v_k, Q, R, V_{oc}, V_p, R_0	Battery latent state (SOC, polarisation voltage), current input (A), terminal-voltage measurement (V), process and measurement noise with covariances Q and R, open-circuit voltage (V), polarisation voltage (V), ohmic resistance (Ω) (Eq. 13)	Specified; not simulated

Symbols	Definition (units)	Status in this paper
$\widehat{D}_t$, p_j, $\hat{n}_{j,t}$, s_t, ν_j, c_t	Forecast instructional demand (W), rated power of one unit of device class j (W), forecast number of active units (–), schedule indicator (0 or 1), demand regressor, calendar covariates (Eq. 14)	Specified; not simulated (loads taken from the generator schedule)
s_t, a_t, m_t, r_t, U_t, B_t, R_t, w_j, $u_{j,t}$, λ_1–λ_3, κ, SOC^*	MDP state and action, relocation flag (0 or 1), reward (–), criticality-weighted unmet energy (weight·kWh), battery-wear penalty (kWh), renewable-utilisation fraction (–), criticality weight (–) and unmet energy (Wh) of class j, reward weights, penalty slope (–), SOC threshold (–) (Eq. 15)	Simulated
γ_d, α, Q^*, Q	Discount factor, learning rate, optimal and estimated action-value functions (Eqs. 16–17)	Simulated (tabular form)
x_j, w_j, L_j, P_{avail}, $\phi(a_t)$, C_{bat}, η_d, Δt	Binary service decision, criticality weight (–), unmet energy (Wh), power available to loads (W), SOC floor set by the action (–), battery energy capacity (Wh), discharge efficiency (–), time step (h) (Eqs. 18–19)	Simulated
y_i, $\hat{y}_i$, N, $\bar{y}$, P_{cap}, $\mathcal{D}$, τ_d, ref	Observed and forecast values (W), sample size, mean (W), array capacity (W), daytime sample set, daytime threshold (fraction of P_{cap}), reference forecaster (Eqs. 20–25)	Computed in Section 5; adopted in the field protocol

Energy Layer: PV Yield, Device Model, Temperature Derating and MPPT

At sizing level the layer follows the standard yield relation:

$$E = P \times \mathrm{PSH} \times \eta \quad (1)$$

where E is daily energy (Wh day^{-1}), P the rated array power (W), PSH peak sun hours (h day^{-1}) and η system efficiency. For simulation-based design the module is described by the five-parameter single-diode model (Song et al., 2021), written for a module of N_s series cells:

$$I = I_{ph} - I_0\left[\exp\left(\frac{V+IR_s}{N_s\, n\, V_T}\right) - 1\right] - \frac{V+IR_s}{R_{sh}} \quad (2)$$

$$V_T = \frac{k\,T}{q} \quad (3)$$

where the symbols are defined in Table 3. Parameterising Equations (2) and (3) from the manufacturer datasheet of the deployed module lets irradiance and temperature sensors feed a physically grounded estimate. These two equations are specified but are not simulated in Section 5. For the forecasting and simulation components, a temperature-corrected power relation of the form used in the PV-performance literature (Skoplaki & Palyvos, 2009) is adopted:

$$P_{pv} = P_{stc}\,\frac{G}{G_{stc}}\left[1 + \gamma(T_c - T_{ref})\right]\eta_{sys}, \quad T_c = T_a + k_G\, G \quad (4)$$

with γ negative for crystalline silicon. Equation (4) is the generation model used in Section 5. The charge controller tracks the maximum power point using a hybrid perturb-and-observe and neural-network scheme of the kind classified by Worku et al. (2023), governed by the stationarity condition on the power–voltage curve:

$$\frac{dP}{dV} = I + V\,\frac{dI}{dV} = 0 \quad (5)$$

The operating voltage is perturbed toward the point where dP/dV vanishes. MPPT is specified but not simulated: it operates inside the charge controller, is not one of the learning components exercised in Section 5, and its effect is subsumed in η_{sys}.

AI Decision Layer

Generation Forecasting

Short-horizon generation is forecast with an LSTM (Hochreiter & Schmidhuber, 1997) in the gated formulation with a forget gate (Gers et al., 2000). With input vector x_t (lagged PV power, clear-sky power, hour-of-day encodings and, in deployment, temperature, humidity and cloud indicators) and previous hidden state h_{t-1}:

$$f_t = \sigma\big(W_f\,[h_{t-1}, x_t] + b_f\big) \quad (6)$$

$$i_t = \sigma(W_i\,[h_{t-1}, x_t] + b_i) \quad (7)$$

$$\tilde{c}_t = \tanh(W_c\,[h_{t-1}, x_t] + b_c) \quad (8)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \quad (9)$$

$$o_t = \sigma(W_o\,[h_{t-1}, x_t] + b_o), \quad h_t = o_t \odot \tanh(c_t) \quad (10)$$

$$\hat{E}_{t+1} = g(h_t; \theta) \quad (11)$$

Equation (11) is a dense read-out of the final hidden state. A gradient-boosted tree model (Chen & Guestrin, 2016) and smart persistence are retained as comparison baselines.

Battery-State Estimation

The baseline estimate is Coulomb counting. With current positive when charging, C_n in ampere-hours and τ in seconds,

$$\mathrm{SOC}(t) = \mathrm{SOC}(0) + \frac{1}{3600\,C_n}\int_0^t \eta_c\; i(\tau)\, d\tau \quad (12)$$

so that SOC is dimensionless and rises while charging, with η_c below one when charging and equal to one when discharging. Because open-loop counting drifts under sensor noise and unknown initial conditions, AI-SEHOS corrects it with a Kalman-type filter on a discrete equivalent-circuit model whose measurement equation is nonlinear in SOC (Plett, 2004; Hu et al., 2022):

$$\begin{aligned} x_k &= A\,x_{k-1} + B\,u_k + w_k, & w_k &\sim \mathcal{N}(0, Q) \\ z_k &= V_{oc}(\mathrm{SOC}_k) - V_{p,k} - R_0\,i_k + v_k, & v_k &\sim \mathcal{N}(0, R) \end{aligned} \quad (13)$$

Here x_k contains SOC and the polarisation voltage V_p, u_k is the measured current and z_k the terminal voltage. Because $V_{oc}(\cdot)$ is a nonlinear function of SOC, the filter is implemented as an extended or unscented filter (Julier & Uhlmann, 2004), with sigma-point propagation in the unscented case and, in the adaptive variant, on-line updating of R from the innovation sequence. The covariances Q and R are site- and cell-specific and must be tuned on telemetry. The corrected SOC, rather than raw counting, enters the dispatch state. The Kalman-type estimator (Equation 13) is specified but not simulated in Section 5, where SOC is updated by noise-free energy bookkeeping equivalent to Equation (12).

Educational Demand Forecasting

Instructional demand is formalised as a forecast of the power requested by each device class j during scheduled sessions:

$$\widehat{D}_t = s_t \sum_j p_j \, \hat{n}_{j,t}, \quad \hat{n}_{j,t} = \nu_j(\mathrm{att}_{t-1}, \ldots, \mathrm{att}_{t-L};\, c_t, m_t) \quad (14)$$

where s_t indicates a scheduled session, p_j is the rated power of one unit of class j (for example a tablet), and $\hat{n}_{j,t}$ is the forecast number of active units from recent attendance att, calendar covariates c_t and the relocation flag m_t. A light regressor, such as a Poisson or gradient-boosted model, is sufficient for the small samples expected at pilot stage. This is the point at which educational information enters the energy problem. It is specified but not simulated: in Section 5 the scheduled loads are generated directly.

Reinforcement-Learning Dispatch

Following the MDP formulation of Alabdullah and Abido (2022) and the comparative review of Mohammadi et al. (2024), dispatch is a sequential decision problem with state $s_t = \left(\mathrm{SOC}_t, \hat{E}_{gen,t}, \widehat{D}_t, \mathrm{hour}, m_t\right)$ and a discrete action a_t chosen from four options that set the SOC floor $\phi(a_t)$ below which non-critical loads may not draw on the battery: charge-priority ($\phi = 0.60$), balanced (0.40), discharge-priority (0.10) and curtail (non-critical loads shed; floor 0.10 for critical loads). The server or connectivity load is always allowed to draw down to the physical minimum of 0.10. The reward balances instructional continuity, battery wear and renewable utilisation:

$$r_t = -\lambda_1 U_t - \lambda_2 B_t + \lambda_3 R_t \quad (15)$$

where $U_t = \sum_j w_j \, u_{j,t}$ is criticality-weighted unmet energy (weight·kWh), $B_t = \Theta_t[1 + \kappa \max(0, \mathrm{SOC}^* - \mathrm{SOC}_{t+1})]$ is a depth-dependent degradation penalty in which Θ_t is battery throughput (kWh), R_t is the fraction of available PV energy used or stored (counted only in hours with PV output), and λ_1–λ_3 are reward weights. The optimal action-value function satisfies the Bellman optimality equation, and tabular Q-learning (Watkins & Dayan, 1992) approximates it by the temporal-difference update:

$$Q^*(s,a) = \mathbb{E}\left[r_t + \gamma_d \max_{a'} Q^*(s_{t+1}, a') \mid s_t = s, a_t = a\right] \quad (16)$$

$$Q(s,a) \leftarrow Q(s,a) + \alpha\left[r_t + \gamma_d \max_{a'} Q(s_{t+1}, a') - Q(s,a)\right] \quad (17)$$

At deployment scale the continuous state calls for a deep Q-network (Mnih et al., 2015) or a policy-gradient method such as DDPG (Lillicrap et al., 2016); Section 5.3 uses the tabular form because the illustrative state space is small, and no claim is made about the deep variants.

Constrained Load Allocation

Given the power released by the dispatch action, the system decides which loads to serve. Let load unit j have rated power p_j (W) and criticality weight w_j, and let $x_j \in \{0,1\}$ indicate that it is served during the interval. The allocation problem is:

$$\min_{x_j \in \{0,1\}} \sum_j w_j \, L_j, \quad L_j = p_j(1 - x_j)\Delta t, \quad \text{s.t. } \sum_j p_j \, x_j \leq P_{avail}(t) \quad (18)$$

$$P_{avail}(t) = P_{pv}(t) + \frac{\eta_d \, C_{bat} \max[0, \mathrm{SOC}(t) - \phi(a_t)]}{\Delta t} \quad (19)$$

where C_{bat} is battery energy capacity in Wh, η_d the discharge efficiency, $\Delta t = 1$ h the time step and $\phi(a_t)$ the SOC floor implied by the dispatch action. The second term of Equation (19) is energy (Wh) divided by time (h) and therefore has units of power (W), consistent with P_{pv} and with the power constraint in Equation (18); L_j is unmet energy in Wh. Critical loads use the physical floor $\phi = 0.10$, so Equation (19) is evaluated twice per interval, once with $\phi = 0.10$ for the server and once with $\phi(a_t)$ for the remaining loads. This links Equation (18) to the dispatch decision explicitly.

How the problem is solved. When loads are divisible and the penalty is proportional to unmet energy, serving loads in decreasing order of w_j is optimal, because the problem is a fractional knapsack whose value density equals w_j; the allocation then depends only on the rank order of the weights. For indivisible units, as in this paper, the problem is a 0–1 knapsack, which is NP-hard in general. The implementation therefore uses a greedy rule: classes are visited in decreasing order of weight and as many whole units of each class as fit within the remaining power are served, so that a unit that does not fit is skipped and smaller units of a lower class can still be served. Because every unit in the scenario is a multiple of 5 W and at most 39 units are scheduled, the exact optimum can be computed by dynamic programming, and Section 5.4 compares the greedy rule with that optimum for every hour in which the power constraint binds. A hard constraint can additionally reserve energy for the server, which is a single point of failure for the whole session. Criticality weights are elicited, not learned: an ordinal default of 5 (server and connectivity), 3 (learner tablets) and 1 (lighting and fans) is used for illustration, to be replaced by weights elicited from teachers and community stakeholders, for example by pairwise comparison.

Educational Service Layer and Logic Model

The Educational Service Layer consumes the allocated energy to run a locally hosted learning-management system, offline content and learner devices. To keep the energy–education link falsifiable, the framework adopts an explicit logic model. Energy availability determines the availability of the server and devices during scheduled sessions (service availability). Service availability determines whether sessions can be held as planned (session continuity). Session continuity, together with acceptance by teachers and caregivers (Davis, 1989; Rogers, 2003), is expected to influence attendance and lesson completion. The framework does not model adaptive-learning algorithms; personalisation, if added, would run on the local content server and enter the energy problem only through Equation (14). Section 6 turns each link into a testable hypothesis.

RESULTS AND DISCUSSION

Illustrative Synthetic Simulations

Design, Settings and Reproducibility

Because no site data exist, the components were tested on a documented synthetic generator. The generator is a transparent test bed, not a model of the real region, and all quantities below are synthetic. It produces hourly PV power from clear-sky irradiance computed from solar geometry at a representative latitude of 12° N, multiplied by a clear-sky index drawn from four weather regimes (clear, harmattan haze, partly cloudy, overcast). Regimes change daily according to a Markov rule with season-dependent frequencies: with probability 0.55 the previous day's regime is retained, and otherwise a new regime is drawn from the seasonal frequencies listed in Table 4. Within a day the clear-sky index follows a first-order autoregressive process around the regime mean, and the temperature derating of Equation (4) is applied. A relocation day occurs with probability 0.08 and removes class load for the day and PV output during hours 8 to 13. Loads are a 40 W server, learner tablets (5 W each, with daily attendance drawn from a Poisson distribution with mean 20, truncated to 8–35) and three 10 W lamps, scheduled for 08:00–14:00 and for two evening hours (18:00–20:00, tablets halved). Mean daily PV output is 1.31 kWh in the dispatch traces (1.38 kWh over the 3.5-year forecasting series, which contains no relocation days), against a mean daily demand of

1.18 kWh (1.26 kWh on a day without relocation), so scarcity arises mainly in multi-day overcast spells. Evaluation uses data never seen in training: chronological splits for forecasting, and trace-disjoint sets for dispatch.

Two generator configurations were run. In configuration A the clear-sky index is a latent autoregressive process only. This was the first configuration run, and under it the learned forecasters did not outperform smart persistence (Section 5.2). Configuration B, the main configuration, adds two features that are physically motivated for the Sahel: hour-to-hour (white) cloud-flicker variability and a diurnal cloud cycle with morning dust loading in the haze regime and afternoon convective cloud in the partly-cloudy regime. B was added to test the sensitivity of forecast skill to unpredictable high-frequency variability and diurnal structure; no other variants were tried, and both configurations are reported. All simulation code, random seeds, generated test predictions and raw outputs are supplied as Supplementary Material (see Data and code availability), and every number in Tables 5–7 and Figures 4–6 is produced by that code. Table 4 lists the settings.

Table 4: Synthetic Scenario, Generator Parameters and Experimental Settings

Item	Setting
Array and system	$P_{stc} = 400$ W (also used as P_{cap}), $\eta_{sys} = 0.80$, $\gamma = -0.004$ °C^{-1}, $T_c = T_a + 0.03\,G$; battery 1,200 Wh, $\eta_c = \eta_d = 0.95$, SOC limits 0.10–1.00, start SOC 0.60; $\Delta t = 1$ h
Generator: weather regimes	Regime mean clear-sky index (clear, haze, partly cloudy, overcast) = 0.94, 0.74, 0.62, 0.22; standard deviation of the autoregressive component = 0.03, 0.05, 0.18, 0.07; hourly autoregressive coefficient 0.90; index clipped to 0.05–1.05. Seasonal regime frequencies (clear, haze, partly cloudy, overcast): harmattan, December–February, 0.50/0.40/0.09/0.01; dry, March–May, 0.55/0.15/0.27/0.03; rainy, June–September, 0.10/0.00/0.45/0.45; post-rain, October–November, 0.45/0.10/0.38/0.07. Ambient temperature = seasonal and diurnal climatology plus Gaussian noise (SD 0.8 °C)
Generator: configuration B additions	White noise on the clear-sky index (SD 0.03, 0.04, 0.07, 0.04 by regime); haze-regime mean lowered by 0.10 in the morning (Gaussian bump centred 08:30, SD 2 h); partly-cloudy mean lowered by 0.12 × seasonal factor (0.10, 0.40, 1.00, 0.60) in the afternoon (bump centred 15:00, SD 2.2 h); clear-regime mean lowered by 0.02 in the afternoon
Relocation and loads	Relocation probability 0.08 per day (class load removed; PV removed 08:00–13:59). Server 40 W; tablets 5 W each, attendance Poisson(20) truncated to 8–35, halved in the evening; three 10 W lamps; schedule 08:00–14:00 and 18:00–20:00
Forecasting data	1,277 days of hourly data without relocation events (30,648 h): first 730 days training (17,496 samples after the 24-h look-back), next 365 days validation (8,760; early stopping), final 182 days test (4,368; 1 January–1 July of the fourth year, covering the harmattan, dry and early rainy seasons). Series seed 42
Forecasting models	Persistence; smart persistence (clear-sky power × last observed clear-sky index); simple exponential smoothing ($\alpha = 0.4$); gradient boosting (250 trees, depth 4, learning rate 0.1, 12 lag, clear-sky and calendar features); LSTM (1 layer, 24 units, 24-h look-back, 5 inputs per step, Adam (Kingma & Ba, 2015) with learning rate 0.01, gradient-norm clipping at 1, batch 128, 14 epochs, best validation epoch kept; implemented in NumPy; five random seeds 0–4, seed 0 reported as the headline model)
Dispatch learning	Tabular Q-learning, state = (SOC bin [6], hour [24], previous-day yield bin [3], relocation flag [2]), 4 actions; $\alpha = 0.1$, $\gamma_d = 0.95$; ϵ decayed 1.0→0.05 over 70% of 2,500 episodes; each episode a random 21-day trace from 60 training traces (seeds 1000–1059); 5 independent agents (seeds 11–15); reward weights $\lambda_1 = 5$, $\lambda_2 = 1$, $\lambda_3 = 2$, $\kappa = 2$, $\mathrm{SOC}^* = 0.5$; learning curve evaluated on 10 validation traces (seeds 2000–2009)
Dispatch evaluation	30 held-out 21-day traces (seeds 3000–3029, disjoint from training and validation); greedy policy; the five agent policies are averaged per trace; results are means with 95% t-confidence intervals across the 30 traces; paired differences are computed trace by trace
Baselines	Fixed charge-priority; fixed balanced; fixed discharge-priority; SOC-threshold rule (charge-priority if SOC < 0.5, otherwise discharge-priority)
Load allocation	Unit devices (server 40 W, tablet 5 W, lamp 10 W); weights (5, 3, 1) unless stated; battery policy fixed at discharge-priority; same 30 held-out traces; greedy rule compared with exact 0–1 knapsack (dynamic programming)
Uncertainty	RMSE intervals: percentile bootstrap over the 182 test days with 1,000 resamples (Efron & Tibshirani, 1993), seed 7; the paired bootstrap uses the same resampled days for both models

Forecasting

Table 5 reports one-hour-ahead performance on the held-out 182 days. Errors are normalised by array capacity (P_{cap} = 400 W; nRMSE) because percentage errors are undefined when output is zero at night; MAPE is computed over daytime hours only, defined as hours when observed output exceeds 10% of capacity (Equation 23; Hyndman & Koehler, 2006). Skill is measured against smart persistence (Equation 25). Panel A gives the main configuration (B) and panel B the sensitivity configuration (A).

Table 5. Held-Out One-Hour-Ahead Forecast Accuracy on Synthetic Data (n = 4,368 Hourly Test Points)

Model	MAE (W)	RMSE (W) [95% CI]	nRMSE (%)	Daytime MAPE (%)	R^2	Skill vs smart persistence
Panel A. Configuration B (main): AR(1) clear-sky index with white noise and diurnal cloud cycle						
Persistence	18.3	28.9 [28.1, 29.6]	7.23	31.9	0.881	−1.702
Smart persistence	4.8	10.7 [9.7, 11.8]	2.67	9.2	0.984	0 (reference)
Exponential smoothing	38.4	54.0 [52.3, 55.6]	13.51	50.6	0.585	−4.050
Gradient boosting	4.8	10.2 [9.2, 11.2]	2.56	8.9	0.985	+0.043
LSTM	5.4	10.1 [9.2, 11.0]	2.52	8.9	0.986	+0.056
Panel B. Configuration A (sensitivity): AR(1) clear-sky index only						
Persistence	17.8	28.1 [27.3, 28.9]	7.03	30.3	0.891	−3.500
Smart persistence	2.5	6.3 [5.5, 6.9]	1.56	4.6	0.995	0 (reference)
Exponential smoothing	38.9	54.7 [53.0, 56.4]	13.68	50.4	0.589	−7.753
Gradient boosting	2.8	6.5 [5.8, 7.2]	1.63	5.1	0.994	−0.041
LSTM	3.5	6.5 [5.9, 7.2]	1.63	5.3	0.994	−0.045

Note: RMSE Values Show 95% Percentile-Bootstrap Intervals over the 182 Test Days in Brackets (1,000 Resamples). Skill = 1 − RMSE/RMSE of Smart Persistence. Models Trained on the First Two Years; the LSTM Row is the Seed-0 Model

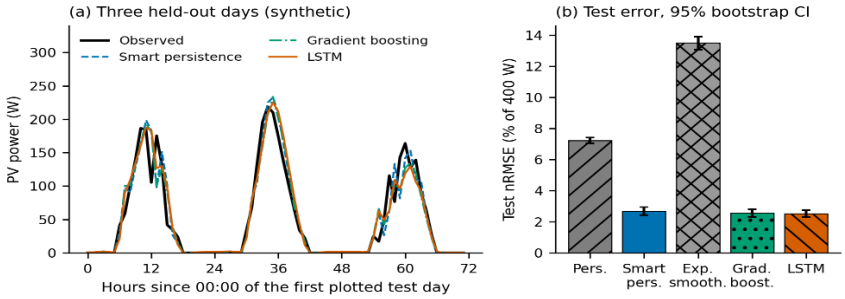


Figure 4: One-Hour-Ahead Forecasts Over a Held-out Three-Day Window (Left) and Test nRMSE with 95% Bootstrap Intervals (right) for Configuration B, Synthetic Data. The Window was Chosen Automatically as the Three-day Window with the Largest Smart-persistence error. All Models and Values Correspond to Table 5, Panel A

In configuration B the LSTM and the gradient-boosting model both achieve a lower RMSE than smart persistence, but the margin is small: skill is 0.056 for the LSTM and 0.043 for gradient boosting, that is, reductions in RMSE of about 6% and 4%. The paired bootstrap places the LSTM skill at 0.056 (95% interval 0.033 to 0.079) and the gradient-boosting skill at 0.043 (0.021 to 0.062), so the improvement over the naive reference is statistically detectable on this series but modest. The difference between the two learned models is -0.14 W in RMSE (95% interval -0.38 to +0.10 W), which includes zero, and the five LSTM seeds span 10.07–10.22 W; the data therefore give no basis for ranking the two models. Errors are concentrated in the partly cloudy regime, where RMSE is about 16 W for the LSTM against about 6 W in the clear regime. Raw persistence and exponential smoothing perform poorly because they ignore the diurnal cycle, which is why smart persistence is the appropriate naive reference.

The result is sensitive to the generator. In configuration A, whose clear-sky index is a purely autoregressive process, smart persistence is close to optimal and both learned models are slightly worse than it (skill -0.045 and -0.041, with paired intervals below zero), so no gain from learning is found. The gain in configuration B arises because the added white noise is unpredictable from the last observation while remaining partly predictable by averaging over recent hours, and because the diurnal cloud cycle can be learned from the hour of day. Real irradiance has cloud dynamics that neither generator reproduces, so the size and even the sign of the gain should not be extrapolated; the experiment shows that forecast skill over smart persistence can be measured with held-out data and that it depends on the structure of the series.

Dispatch

Table 6 and Figure 5 report dispatch performance on the 30 held-out traces. Training, validation and evaluation traces are disjoint, and the agent was trained from five different random seeds.

Table 6: Dispatch Results on 30 Held-out 21-day Synthetic Traces (Means ± 95% Confidence-Interval Half-Widths Across Traces)

Policy	Cumulative reward	Weighted unmet load (weight·kWh)	Unmet energy (%)	Server available in school hours (%)	PV utilisation (%)
Fixed: charge-priority	344.6 ± 14.6	7.53 ± 3.22	13.2 ± 5.3	100.0 ± 0.0	82.6 ± 3.5
Fixed: balanced	354.0 ± 15.1	6.73 ± 3.15	11.7 ± 5.2	100.0 ± 0.0	83.9 ± 3.6
Fixed: discharge-priority	358.9 ± 15.7	6.38 ± 3.32	10.2 ± 5.1	98.0 ± 1.4	84.9 ± 3.6
SOC-threshold rule	351.6 ± 15.1	7.04 ± 3.20	12.0 ± 5.2	100.0 ± 0.0	83.6 ± 3.6
Tabular Q-learning (5 seeds)	354.3 ± 15.1	6.79 ± 3.28	11.1 ± 5.1	98.5 ± 1.3	84.4 ± 3.6

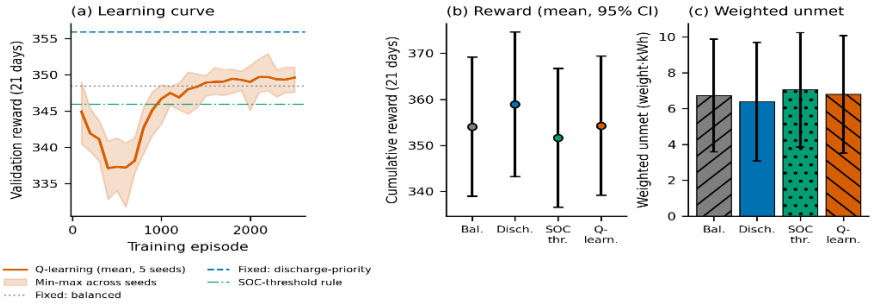


Figure 5: Learning Curve of the Tabular Agent on Validation Traces (Left; Mean of Five Seeds, Band = Min–Max Across Seeds; Horizontal Lines = Rule Baselines on the Same Traces), Cumulative Reward on Held-Out Traces (Centre; Markers = Means, Bars = 95% Confidence Intervals) and Weighted Unmet Load (Right), Synthetic Data

Paired differences across traces show that the Q-learning policy is indistinguishable from the fixed balanced rule in reward (+0.21 ± 1.96), clearly worse than the fixed discharge-priority rule (−4.65 ± 1.62), and better than the fixed charge-priority rule (+9.62 ± 3.18) and, marginally, the SOC-threshold rule (+2.63 ± 2.17). Differences in weighted unmet load follow the same pattern (+0.06 ± 0.26, +0.41 ± 0.15, −0.74 ± 0.35 and −0.26 ± 0.27 weight·kWh, respectively). Seed-to-seed variation in mean test reward was small (353.7–355.1), and the learned policy uses all four actions (charge-priority 27%, balanced 28%, discharge-priority 29%, curtail 16% of hours), so it is not simply reproducing one rule. The learned policy therefore matches the middle of the rule-based comparators but does not outperform the best of them, fixed discharge-priority, in this setting; on the validation traces the learning curve also remains below that rule after 2,500 episodes. The reason is plausible: with a small battery, deterministic schedules and an engineered reward, a single well-chosen rule is hard to beat, and a tabular state offers little extra information. The results neither support nor refute the larger gains reported for deep RL in microgrid studies (Alabdullah & Abido, 2022; Mohammadi et al., 2024), which used different systems and function approximation. The case for RL in this framework is therefore an open empirical question, and any field study should include rule-based controllers as baselines (Section 6).

Load Prioritisation

Table 7 and Figure 6 compare allocation rules under the fixed discharge-priority battery policy. The naive rule serves loads in installation order (lighting, tablets, server); criticality-aware rules serve in descending weight. Three weight sets preserve the rank server > tablets > lighting but differ widely in magnitude, and one reverses the ranking.

Table 7: Load Allocation on 30 Held-out Traces: Unmet Energy by Class, Server Availability and Weighted Cost (Means; Weighted Cost in Weight·kWh Over 21 Days, ± 95% Confidence-Interval Half-Width)

Allocation rule	Server unmet (%)	Tablets unmet (%)	Lighting unmet (%)	Server available in school hours (%)	Cost under own weights	Cost under (5, 3, 1)
Naive fixed order	17.6	9.7	1.6	81.4	9.66 ± 4.81	9.66 ± 4.81
Weights (5, 3, 1)	3.5	10.7	17.4	98.0	6.38 ± 3.32	6.38 ± 3.32
Weights (10, 2, 1)	3.5	10.7	17.4	98.0	5.99 ± 3.10	6.38 ± 3.32
Weights (3.0, 2.9, 2.8)	3.5	10.7	17.4	98.0	7.26 ± 3.70	6.38 ± 3.32
Reversed (1, 3, 5)	17.6	9.7	1.6	81.4	5.54 ± 2.91	9.66 ± 4.81

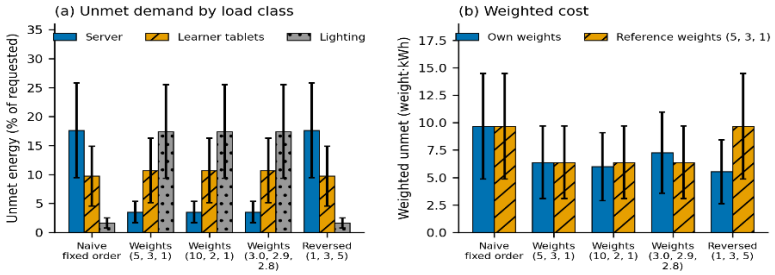


Figure 6: Unmet Demand by Load Class (Left; Bars = 95% Confidence Intervals Across Traces) and Weighted Cost Under Each Rule's Own Weights and Under the Reference Weights (Right), Synthetic Data

Three results follow. First, criticality-aware allocation raises server availability during school hours from 81.4% to 98.0% and lowers cost under the reference weights from 9.66 to 6.38, a reduction of about 34% (paired per-trace reduction 34% ± 16%), at the price of more unmet lighting (1.6% to 17.4% of requested lighting energy). Second, all three rank-preserving weight sets give identical allocations, whereas the rank-reversed set reproduces the naive allocation, because the naive order is the reverse of the criticality ranking. Allocation therefore depends only on rank. Third, the size of the weighted-cost figure does depend on magnitude: the same allocation costs 5.99, 6.38 or 7.26 under the three rank-preserving sets, and the reversed set appears to "win" (5.54) only because it is scored with its own weights, while under the reference weights it costs 9.66. A weighted cost is thus not

comparable across weight sets, and the 34% figure applies to the reference weights only.

Practical interpretation. The trade-off is between the server, on which every learner device depends, and lighting, which is a low-power convenience: lighting accounts for 30 W of a scheduled load of about 1.26 kWh per day. Under criticality-aware allocation the server stays available in about 98% of school hours, so classes can run, while lamps are more often unlit when the battery is depleted (about 17% of requested lighting energy is unmet, against 2% under the naive rule). Whether this is acceptable is a question for the community: if evening study depends on light, a small stand-alone solar lantern or a protected lighting circuit could be provided so that lamps do not compete with the server, and the weights should be elicited accordingly (Section 4.3.5).

Because Equation (18) is solved here by a greedy sorting rule, the experiment is a demonstration of its consequences, not a test of an algorithm. To check that the greedy rule is not itself a source of loss, it was compared with the exact 0–1 knapsack optimum by dynamic programming in every scheduled hour of the 30 held-out traces under fixed discharge-priority: of 4,728 scheduled hours, the power constraint was binding in 844, and the greedy allocation attained the exact optimum in 844 of them (mean gap 0.0%). This is a property of the small unit sizes in this scenario, not a general guarantee. The experiment shows what the weights change, namely who loses power, and not whether the weights match what teachers value.

What the Simulations Do and Do Not Show

They show that the specified components run end to end, that forecasting skill over a proper naive reference can be measured with held-out data and depends on the structure of the series, that a simple learned dispatch policy can match but not exceed rule-based controllers, and that load allocation redistributes shortfalls according to the ranking of criticality weights. They do not show field performance, LSTM or RL superiority on real data, optimal weights, battery-ageing benefits, cost-effectiveness or educational impact. The Kalman-type SOC estimator and MPPT were not simulated. The generator's weather regimes are assumptions, not measurements.

Evaluation Protocol for Field Validation

The protocol below is prospective: no part of it has been executed. Performance would be assessed against four coordinated families of indicators (Table 8), each with a hypothesis and a pre-stated threshold, set before data collection so that they can be falsified, within the pilot design described in Section 3.2. Forecast accuracy uses:

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}|\hat{y}_i - y_i| \quad (20)$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2} \quad (21)$$

$$\mathrm{nRMSE} = 100 \cdot \frac{\mathrm{RMSE}}{P_{cap}} \quad (22)$$

$$\mathrm{MAPE} = \frac{100}{|\mathcal{D}|}\sum_{i\in\mathcal{D}}\left|\frac{y_i-\hat{y}_i}{y_i}\right|, \quad \mathcal{D} = \{i: y_i > \tau_d P_{cap}\} \quad (23)$$

$$R^2 = 1 - \frac{\sum_i (y_i-\hat{y}_i)^2}{\sum_i (y_i-\bar{y})^2} \quad (24)$$

$$\mathrm{Skill} = 1 - \frac{\mathrm{RMSE}}{\mathrm{RMSE}_{ref}} \quad (25)$$

Equation (23) restricts MAPE to daytime hours because it is undefined at $y_i = 0$ and unstable near zero (Hyndman & Koehler, 2006). Smart persistence is the default reference in Equation (25). Precision, recall and F1 are reported for event-detection sub-tasks such as low-generation alerts.

Table 8: Prospective Evaluation Protocol: Indicators, Hypotheses and Proposed Thresholds (to be Fixed Before Data Collection; Estimation Targets for a Feasibility-Scale Pilot)

Family	Indicators	Hypothesis	Proposed threshold
Energy	Harvested energy; renewable-utilisation ratio; battery round-trip efficiency; SOC error (Eqs. 12–13); unserved-load hours	H1: Field-measured generation is explained by Eq. (4) within tolerance after site calibration.	Daily-energy error ≤ 15% on ≥ 8 weeks of data
AI	MAE, RMSE, nRMSE, daytime MAPE, R², skill (Eqs. 20–25); cumulative reward (Eq. 15)	H2: The LSTM beats smart persistence out of sample. H3: The learned dispatch policy beats the best rule-based baseline.	H2: skill ≥ 0.05 with the paired interval above 0 (the gain in Section 5.2 was about 0.05). H3: weighted unmet load reduced ≥ 10% with non-overlapping intervals; H3 is explicitly open given Section 5.3
Service	Server and device availability in scheduled hours; critical-load unmet energy	H4: Criticality-aware allocation keeps the server available during scheduled sessions.	≥ 95% of scheduled school hours
Educational and adoption	Sessions held as planned; attendance; device-use hours; lesson completion; Technology Acceptance Model constructs (perceived usefulness, ease of use); local custodianship	H5: Service availability is associated with session continuity and attendance. H6: Teachers and caregivers report acceptable usefulness and ease of use.	H5: ≥ 90% of scheduled sessions held; association estimated with confidence intervals. H6: pre-registered scale cut-offs

Reporting the four families jointly, rather than energy-engineering or educational metrics alone, is part of the framework's intended contribution. Unmatched comparisons should be avoided: AI components must be compared with rule-based controllers on the same site data, using the within-site alternating-block design of Section 3.2, and costs recorded alongside (Section 7.5).

Discussion

Contributions

The paper contributes four things, each bounded by the evidence presented. (i) An equation-level specification in which instructional demand and mobility appear explicitly in the forecasting, dispatch and allocation components (Equations 14, 15 and 18), rather than as fixed background load. (ii) A reproducible synthetic test bed and held-out experiments that show how the components behave relative

to simple baselines. (iii) A logic model and evaluation protocol with pre-stated hypotheses (Table 8). (iv) A candid account of what has not been shown. These are design contributions; they do not demonstrate that the framework improves learning or that it outperforms alternatives in the field.

Positioning Against Prior Work

Within the sources examined (Section 2.7), the AI-for-energy literature optimises technical objectives such as yield, curtailment and cost (Ejiyi et al., 2025; Ukoba et al., 2024), and the AI-in-education literature assumes stable power (Alfredo et al., 2024). AI-SEHOS differs by placing the instructional use of energy inside the objective, and by acknowledging mobility. Whether this difference matters in practice depends on how much unmet demand falls on the educationally critical loads, which the allocation experiment shows is a matter of weights, and on whether community-elicited weights differ from the engineering default. The paper's novelty claim is accordingly limited, and a systematic mapping is needed before a stronger one is made.

Limitations

All results are synthetic, from a generator whose weather regimes and load schedules are assumptions, and the forecasting conclusions change with the generator (Section 5.2). The forecasting test period covers about six months of one simulated series, with no multi-year or multi-site variation. The dispatch experiment uses a small tabular agent and a single battery size. The allocation result follows from the ranking by construction. The Kalman-filter component (Equation 13) and MPPT (Equation 5) were specified but not simulated, and the Kalman filter depends on noise covariances that require site-specific tuning. The reviewed literature is purposive (Section 2.7). Reward and criticality weights are placeholders. Non-technical barriers documented in the nomadic-education literature, including curriculum rigidity, teacher shortages and the time needed to build caregiver trust, cannot be resolved by an energy-and-AI framework, which is at best an enabling condition (Aderinoye et al., 2007).

Computational Feasibility

LSTM or CNN forecasting, filter-based SOC estimation and RL dispatch are typically trained on powerful hardware, which conflicts with the offline-first, low-power setting. A split architecture is the plausible response: a frozen forecaster and a small pre-trained dispatch policy (the Q-table in Section 5.3 has 864 states and four actions, about 14 kB at 32-bit precision) run on low-power embedded hardware of the kind surveyed in the tiny-machine-learning literature (Alajlan & Ibrahim, 2022; Ray, 2022), while retraining occurs off-site when connectivity allows. Prototyping should report inference latency and power draw, so that the compute load can be shown to be small relative to the energy it manages. The comparable performance of gradient boosting and the LSTM in Table 5 also suggests that a lighter model may suffice.

Cost Feasibility

Affordability may be the binding constraint on adoption, and this paper does not estimate it. In the simulated scenario, a 400 W array and 1,200 Wh battery serve a demand of about 1.2 kWh per day with some unmet load in overcast spells, which gives a sizing basis but not a price. A field phase should record a bill of materials (array, controller, battery, inverter, sensors, embedded compute, mounting and protective housing), installation and maintenance costs, and compute cost per learner-hour and levelised energy cost against the generator or no-power baseline before any claim of financial sustainability is made. Continental estimates of the cost of powering fixed schools with solar PV (Moner-Girona et al., 2025) are a useful external reference but do not cover mobile sites. Sustainability more broadly, including battery end-of-life handling, local repair capacity and theft risk (Table 2), is likewise a field question.

Transferability

The architecture is not tied to Northern Nigeria, and the continental evidence on school electrification (Moner-Girona et al., 2025) suggests comparable need for mobile or seasonally underserved learners elsewhere in the Sahel and East African pastoral corridors. Any transfer would require re-estimation of the generator parameters, loads, weights and acceptance indicators for the new setting.

CONCLUSION

This paper specifies AI-SEHOS, a four-layer framework that treats instructional demand and mobility as explicit elements of solar energy management for off-grid nomadic learning centres, and exercises its components on a documented, reproducible synthetic test bed with held-out evaluation. Its main contribution is this integrated, equation-level specification together with a pre-stated field-validation protocol; within the sources examined, the combination is not addressed elsewhere, although that statement is limited by the purposive review. The simulations show that learned forecasters give at most a small gain over smart persistence that depends on the generator, that a tabular RL dispatch policy matches but does not exceed the best simple rule, and that criticality-aware allocation redistributes shortages toward the server according to the ranking of weights. The principal limitation is that all evidence is synthetic: field performance, educational impact and cost-effectiveness remain untested.

Future work should first conduct a registered systematic mapping of the literatures listed in Section 2.7, and then validate the forecasting and SOC components against measured irradiance, temperature and battery data from at least one site, benchmarking them against rule-based controllers with the thresholds of Table 8. In parallel, the Kalman-type SOC estimator should be implemented and tuned, deep RL and policy-gradient dispatch tested against the rule baselines on real or higher-fidelity data, and the system prototyped on embedded hardware with latency, power and cost reported. Criticality weights and acceptance indicators should be elicited from teachers, caregivers and programme officers under ethics approval and the Nigeria Data Protection Act, 2023, and educational outcomes should be evaluated alongside energy and AI metrics.

REFERENCES

Aderinoye, R. A., Ojokheta, K. O., & Olojede, A. A. (2007). Integrating mobile learning into nomadic education programmes in Nigeria: Issues and perspectives. *The International Review of Research in Open and Distributed Learning, 8*(2). https://doi.org/10.19173/irrodl.v8i2.347

Alabdullah, M. H., & Abido, M. A. (2022). Microgrid energy management using deep Q-network reinforcement learning. *Alexandria Engineering Journal, 61*(11), 9069–9078. https://doi.org/10.1016/j.aej.2022.02.042

Alajlan, N. N., & Ibrahim, D. M. (2022). TinyML: Enabling of inference deep learning models on ultra-low-power IoT

edge devices for AI applications. *Micromachines, 13*(6), Article 851. https://doi.org/10.3390/mi13060851

Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D., & Martinez-Maldonado, R. (2024). Human-centred learning analytics and AI in education: A systematic literature review. *Computers and Education: Artificial Intelligence, 6*, Article 100215. https://doi.org/10.1016/j.caeai.2024.100215

Bawonda, F. I., Oricha, J. Y., & Sanni, S. O. (2023). Estimating global horizontal irradiance in Nigeria: An empirical modelling perspective. *Journal of Solar Energy Research, 8*(4), 1701–1714. https://doi.org/10.22059/jser.2024.363403.1337

Carabajal, A. T., Orsot, A., Moudio, M. P. E., Haggai, T., Okonkwo, C. J., Jarrard, G. T., III, & Selby, N. S. (2024). Social and economic impact analysis of solar mini-grids in rural Africa: A cohort study from Kenya and Nigeria. *Environmental Research: Infrastructure and Sustainability, 4*(2), Article 025005. https://doi.org/10.1088/2634-4505/ad4ffb

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

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly, 13*(3), 319–340. https://doi.org/10.2307/249008

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Chapman & Hall.

Ejiyi, C. J., Cai, D., Thomas, D., Obiora, S., Osei-Mensah, E., Acen, C., Eze, F. O., Sam, F., Zhang, Q., & Bamisile, O. O. (2025). Comprehensive review of artificial intelligence applications in renewable energy systems: Current implementations and emerging trends. *Journal of Big Data, 12*, Article 169. https://doi.org/10.1186/s40537-025-01178-7

Eldridge, S. M., Lancaster, G. A., Campbell, M. J., Thabane, L., Hopewell, S., Coleman, C. L., & Bond, C. M. (2016). Defining feasibility and pilot studies in preparation for randomised controlled trials: Development of a conceptual framework. *PLOS ONE, 11*(3), Article e0150205. https://doi.org/10.1371/journal.pone.0150205

Gers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural Computation, 12*(10), 2451–2471. https://doi.org/10.1162/089976600300015015

Hennink, M., & Kaiser, B. N. (2022). Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social Science & Medicine, 292*, Article 114523. https://doi.org/10.1016/j.socscimed.2021.114523

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105. https://doi.org/10.2307/25148625

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

Hu, L., Hu, R., Ma, Z., & Jiang, W. (2022). State of charge estimation and evaluation of lithium battery using Kalman filter algorithms. *Materials, 15*(24), Article 8744. https://doi.org/10.3390/ma15248744

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting, 22*(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001

Inman, R. H., Pedro, H. T. C., & Coimbra, C. F. M. (2013). Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science, 39*(6), 535–576. https://doi.org/10.1016/j.pecs.2013.06.002

Julier, S. J., & Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE, 92*(3), 401–422. https://doi.org/10.1109/JPROC.2003.823141

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, 82*(1), 35–45. https://doi.org/10.1115/1.3662552

Kebir, N., Leonard, A., Downey, M., Jones, B., Rabie, K., Bhagavathy, S. M., & Hirmer, S. A. (2023). Second-life battery systems for affordable energy access in Kenyan primary schools. *Scientific Reports, 13*, Article 1374. https://doi.org/10.1038/s41598-023-28377-7

Khouili, O., Hanine, M., Louzazni, M., López Flores, M. Á., García Villena, E., & Ashraf, I. (2025). Evaluating the impact of deep learning approaches on solar and photovoltaic power forecasting: A systematic review. *Energy Strategy Reviews, 59*, Article 101735. https://doi.org/10.1016/j.esr.2025.101735

Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR 2015)*. https://arxiv.org/abs/1412.6980

Koima, J. (2024). School electrification and academic outcomes in rural Kenya. *Journal of Development Economics, 166*, Article 103178. https://doi.org/10.1016/j.jdeveco.2023.103178

Krätli, S., & Dyer, C. (2009). *Mobile pastoralists and education: Strategic options* (Education for Nomads Working Paper 1). International Institute for Environment and Development. https://www.iied.org/sites/default/files/pdfs/migrate/10021IIED.pdf

Lancaster, G. A., Dodd, S., & Williamson, P. R. (2004). Design and analysis of pilot studies: Recommendations for good practice. *Journal of Evaluation in Clinical Practice, 10*(2), 307–312. https://doi.org/10.1111/j..2002.384.doc.x

Leng, F., Tan, C. M., & Pecht, M. (2015). Effect of temperature on the aging rate of Li ion battery operating above room temperature. *Scientific Reports, 5*, Article 12967. https://doi.org/10.1038/srep12967

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations (ICLR 2016)*. https://arxiv.org/abs/1509.02971

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. https://doi.org/10.1038/nature14236

Mohammadi, P., Darabi, R., Shamaghdari, S., & Siano, P. (2024). Comparative analysis of control strategies for microgrid energy management with a focus on reinforcement learning. *IEEE Access*, *12*, 171368–171395. https://doi.org/10.1109/ACCESS.2024.3495032

Moner-Girona, M., Fahl, F., Kakoulaki, G., Kim, D.-H., Maduako, I., Szabó, S., Nhamo, G., Sovacool, B. K., & Weiss, D. J. (2025). Empowering quality education through sustainable and equitable electricity access in African schools. *Joule*, *9*(2), Article 101804. https://doi.org/10.1016/j.joule.2024.12.005

NASA. (n.d.). *Prediction of Worldwide Energy Resources (POWER) data access viewer* [Data set]. https://power.larc.nasa.gov/

Nigeria Data Protection Act, No. 37 of 2023, Federal Republic of Nigeria Official Gazette, Vol. 110, No. 119 (2023). https://ncc.gov.ng/media/801/view

Olaniran, S. O. (2018). Almajiri education: Policy and practice to meet the learning needs of the nomadic population in Nigeria. *International Review of Education*, *64*(1), 111–126. https://doi.org/10.1007/s11159-018-9705-2

Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, *7*(3), 381–388. https://doi.org/10.1109/TII.2011.2158841

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Plett, G. L. (2004). Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Background. *Journal of Power Sources*, *134*(2), 252–261. https://doi.org/10.1016/j.jpowsour.2004.02.031

Ray, P. P. (2022). A review on TinyML: State-of-the-art and prospects. *Journal of King Saud University – Computer and Information Sciences*, *34*(4), 1595–1623. https://doi.org/10.1016/j.jksuci.2021.11.019

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Skoplaki, E., & Palyvos, J. A. (2009). On the temperature dependence of photovoltaic module electrical performance: A review of efficiency/power correlations. *Solar Energy*, *83*(5), 614–624. https://doi.org/10.1016/j.solener.2008.10.008

Song, Z., Fang, K., Sun, X., Liang, Y., Lin, W., Xu, C., Huang, G., & Yu, F. (2021). An effective method to accurately extract the parameters of single diode model of solar cells. *Nanomaterials*, *11*(10), Article 2615. https://doi.org/10.3390/nano11102615

Ukoba, K., Olatunji, K. O., Adeoye, E., Jen, T.-C., & Madyira, D. M. (2024). Optimizing renewable energy systems through artificial intelligence: Review and future prospects. *Energy & Environment*, *35*(7), 3833–3879. https://doi.org/10.1177/0958305X241256293

Umar, S. I., & Maaji, S. M. (2025). From sacred education to street exploitation: The Almajiri crisis in Nigeria as a nexus of public health failures, legal paralysis, and global security risks. *Philosophy, Ethics, and Humanities in Medicine*, *20*(1), Article 23. https://doi.org/10.1186/s13010-025-00191-1

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*(3–4), 279–292. https://doi.org/10.1007/BF00992698

Worku, M. Y., Hassan, M. A., Maraaba, L. S., Shafiullah, M., Elkadeem, M. R., Hossain, M. I., & Abido, M. A. (2023). A comprehensive review of recent maximum power point tracking techniques for photovoltaic systems under partial shading. *Sustainability*, *15*(14), Article 11132. https://doi.org/10.3390/su151411132

World Bank Group. (n.d.). *Global Solar Atlas* [Data set]. https://globalsolaratlas.info/

Yang, D., Kleissl, J., Gueymard, C. A., Pedro, H. T. C., & Coimbra, C. F. M. (2018). History and trends in solar irradiance and PV power forecasting: A preliminary assessment and review using text mining. *Solar Energy*, *168*, 60–101. https://doi.org/10.1016/j.solener.2017.11.023

©2026 This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 International license viewed via https://creativecommons.org/licenses/by/4.0/ which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is cited appropriately.