ISSN: 2096-3246
Volume 58, Issue 02, 2026

# AI-DRIVEN SOLAR POWER OPTIMIZATION AND PRIORITY-BASED LOAD CONTROL SYSTEM WITH VOICE ASSISTANCE

**Dinesh G[1*], Abisha V U[2], Dhanushraj S[3], Ronith Kanna S[4], Mr. S. Bharath[5]**
[1,2,3,4]Department of Electrical and Electronics Engineering, SNS College of Technology, Coimbatore – 641035, Tamil Nadu, India
[5]Professor, Department of Electrical and Electronics Engineering, SNS College of Technology, Coimbatore – 641035, Tamil Nadu, India
***Corresponding Author:** Dinesh G, Department of Electrical and Electronics Engineering, SNS College of Technology, Coimbatore – 641035, Tamil Nadu, India

**Abstract**
The escalating global demand for clean energy has underscored the necessity for intelligent management frameworks that can simultaneously optimize photovoltaic generation, battery storage, and load consumption in dynamic environments. Conventional solar photovoltaic systems operate predominantly in passive modes; they harvest, store, and distribute energy without predictive foresight or autonomous adaptation to supply variability. Consequently, these systems frequently exhibit inefficient power utilization, unplanned battery deep-discharge events, and limited user interactivity. This paper proposes an integrated Smart Solar Energy Management System that synergistically combines Internet of Things (IoT) sensing, hybrid deep learning-based forecasting, priority-driven load scheduling, and natural language voice assistance. An ESP32 microcontroller serves as the embedded edge node, orchestrating real-time acquisition of voltage, current, and environmental data through calibrated analog and digital sensors. The acquired telemetry is transmitted to a cloud-native dashboard for remote visualization and historical logging. At the analytics core, a hybrid neural architecture fusing Gated Recurrent Unit (GRU) and Long Short-Term Memory (LSTM) layers is trained on multivariate time-series data to forecast short-term solar generation trajectories. The predictive intelligence is subsequently fused with a three-tier priority load-control algorithm that dynamically connects or disconnects appliances based on battery state-of-charge thresholds and anticipated energy availability. Furthermore, a voice-assistance module is integrated to enable hands-free inquiry and actuation, thereby improving accessibility and user experience. Extensive field evaluation over a continuous 30-day deployment demonstrates that the proposed hybrid predictor achieves a root-mean-square error of 0.11 kWh and a coefficient of determination of 0.95, outperforming standalone LSTM, GRU, and classical multilayer perceptron baselines. The intelligent load-management strategy reduced deep-discharge incidents by approximately 75% and improved overall energy utilization efficiency by 22%. The voice interface attained 95% recognition accuracy with an average end-to-end latency below 1.5 seconds. These results confirm that the confluence of IoT, hybrid deep learning, and voice-driven automation provides a robust, scalable paradigm for next-generation residential solar energy systems.

**Keywords:** Solar energy management, hybrid deep learning, GRU-LSTM, ESP32, Internet of Things, priority-based load control, voice automation, energy forecasting.

## I. INTRODUCTION

The worldwide transition toward decarbonized electricity networks has positioned solar photovoltaic (PV) technology as a cornerstone of distributed renewable generation. In contrast to centralized fossil-fuel plants, rooftop and standalone solar installations offer modularity, reduced transmission losses, and near-zero operational carbon emissions. According to recent industry projections, global cumulative PV capacity is expected to exceed terawatt-scale deployment within the decade, driven by policy incentives and declining module costs. Despite this remarkable growth, the inherent intermittency of solar irradiance—stemming from diurnal cycles, seasonal variations, cloud transients, and atmospheric unpredictability—poses substantial challenges to power quality, system stability, and end-use reliability. The stochastic nature of1 PV generation necessitates buffering mechanisms, typically electrochemical storage, alongside intelligent dispatch strategies that align energy supply with consumer demand. Without such intelligence, standalone solar systems remain vulnerable to chronic under-performance, premature battery degradation, and service interruptions during low-insolation periods.

Traditional standalone solar architectures generally comprise PV panels, lead-acid or lithium-ion batteries, charge controllers, and manually operated AC or DC loads. While these configurations are functionally adequate for basic electrification, they operate in an open-loop fashion. Charge controllers regulate voltage and current using fixed set-points, and load management is either entirely manual or restricted to simplistic low-voltage disconnect thresholds. As a result, critical and non-critical loads are treated with equal importance, frequently causing deep discharge events when user consumption outpaces generation. Deep discharges below 20% state-of-charge (SOC) are particularly deleterious for lithium-ion cells and can permanently reduce available capacity in lead-acid chemistry. Moreover, users lack visibility into future generation profiles, preventing proactive demand-side adjustments. Consequently, conventional systems sacrifice battery longevity and fail to maximize the self-consumption fraction of locally generated renewable energy.

Over the past decade, the proliferation of low-cost IoT hardware and ubiquitous wireless connectivity has enabled a new class of solar monitoring platforms. Contemporary systems leverage microcontrollers, wireless sensor networks, and cloud gateways to stream real-time telemetry—voltage, current, power, and SOC—to remote dashboards and mobile applications. These platforms deliver valuable operational transparency, allowing users and technicians to diagnose faults, track historical performance, and receive threshold-based alerts. Nevertheless, the vast majority of IoT-enabled solar monitors remain fundamentally reactive; they record and display data without synthesizing predictive insights or autonomously orchestrating load behavior. Data-rich yet decision-poor architectures represent an incomplete solution to the energy management problem, leaving users to manually interpret trends and adjust consumption patterns.

Parallel advances in artificial intelligence, particularly deep learning, have created unprecedented opportunities for time-series forecasting in the energy domain. Recurrent neural architectures such as the Long Short-Term Memory (LSTM) network and the Gated Recurrent Unit (GRU) have demonstrated exceptional proficiency in capturing long-range temporal dependencies within non-stationary sequences. LSTM networks, introduced by Hochreiter and Schmidhuber, employ sophisticated gating mechanisms—input, forget, and output gates—to regulate information flow across arbitrary time lags, rendering them well-suited to meteorological and generation forecasting tasks. GRU variants, proposed by Cho et al., consolidate gating logic into update and reset gates, substantially reducing trainable parameters and computational latency while preserving predictive fidelity. Although both architectures have been applied independently to solar irradiance and power forecasting, standalone models frequently exhibit either slow convergence or overfitting when confronted with highly volatile weather regimes. Hybridization strategies that ensemble complementary recurrent structures offer a compelling pathway to enhance robustness, yet their integration within embedded solar management systems remains largely unexplored.

Beyond forecasting and monitoring, the user experience layer of residential energy systems has garnered increasing attention. Voice-driven human-machine interaction, popularized by consumer virtual assistants, presents a natural modality for smart-home and off-grid energy control. Hands-free inquiry of battery status, generation forecasts, and on-demand load actuation is particularly valuable for elderly users, mobility-impaired individuals, and scenarios where physical display access is inconvenient. While isolated implementations of voice-controlled switching have been documented, the fusion of voice assistance with predictive energy forecasting and automated load prioritization has not been comprehensively addressed in the literature.

Motivated by the aforementioned gaps, this paper presents a unified AI-driven solar power optimization and load control system augmented with voice assistance. The specific research contributions of this work are fourfold. First, a scalable IoT sensing and communication architecture based on the ESP32 microcontroller is designed to acquire, process, and transmit high-fidelity PV, battery, and load telemetry to a cloud analytics layer. Second, a novel hybrid GRU-LSTM deep learning model is developed to forecast short-term solar energy generation using multivariate historical and real-time inputs; the ensemble design capitalizes on the long-term memory stability of LSTM and the computational agility of GRU. Third, a dynamic priority-based load scheduling algorithm is formulated, coupling SOC thresholds with predictive generation availability to autonomously regulate three tiers of consumer appliances, thereby minimizing deep-discharge events and maximizing critical load availability. Fourth, a voice-assistance interface is integrated to deliver intuitive, hands-free system interaction. The proposed framework is validated through a 30-day field deployment, with quantitative performance benchmarks across forecasting accuracy, battery protection efficacy, and voice recognition reliability.

The remainder of this paper is organized as follows. Section II reviews pertinent literature in IoT solar monitoring, deep learning forecasting, and voice control. Section III elaborates the system architecture, hybrid prediction model, priority load-control strategy, and voice interface design. Section IV describes hardware prototyping, dataset curation, and model training procedures. Section V presents experimental results and comparative analysis. Finally, Section VI concludes the paper with a discussion of limitations and future research directions.

## II. LITERATURE REVIEW

### A. IoT-Based Solar Monitoring and Control

The integration of IoT technologies into renewable energy infrastructures has transformed passive generation assets into connected, observable cyber-physical systems. Kekre and Gawre [4] demonstrated an early IoT-enabled remote monitoring system for solar PV installations, utilizing wireless sensors and web dashboards to visualize panel voltage and current. Their work established the feasibility of low-cost telemetry but did not incorporate automated actuation or predictive analytics. Iqbal et al. [5] proposed a data-analytic framework for the design and operation of residential PV systems, emphasizing statistical performance indicators; however, the control loop remained open, with no automated load management. More recent investigations have explored cloud-native architectures and MQTT-based messaging for aggregated solar farms, achieving high data throughput and multi-node scalability. Yet, as highlighted in contemporary reviews, most deployed IoT solar platforms function as decision-support tools rather than autonomous energy optimizers, leaving a significant opportunity to close the sense-predict-act loop at the network edge.

### B. Deep Learning for Solar Energy Forecasting

Accurate forecasting is prerequisite to effective solar energy management. Early efforts relied on physical models and statistical regressions, which require domain-specific assumptions and often fail to generalize across divergent climatic zones. Machine learning classifiers, including Support Vector Machines and shallow Artificial Neural Networks, improved non-linear mapping but were limited by manual feature engineering and constrained memory horizons. The advent of deep recurrent networks revolutionized time-series prediction by learning hierarchical spatio-temporal representations directly from raw or minimally processed sequences. Siami-Namini et al. [10] conducted a rigorous comparison between autoregressive integrated moving average (ARIMA) and LSTM models, concluding that LSTM consistently outperforms classical statistical methods on non-stationary energy datasets. Voyant et al. [9] provided a comprehensive review of machine learning techniques for solar radiation forecasting, identifying deep learning as the predominant paradigm for intra-hour and intra-day horizons. Despite these advances, existing literature predominantly evaluates standalone LSTM or GRU models in offline simulation environments. The operational deployment of hybrid recurrent architectures within resource-constrained embedded solar systems, and their coupling with real-time control policies, remains conspicuously absent.

### C. Priority-Based Load Management Strategies

Demand-side management in off-grid and hybrid solar systems seeks to align controllable loads with available generation. Priority-based load shedding has been extensively studied in microgrid and utility-scale contexts, where hierarchical disconnection schemes preserve grid stability during supply shortfalls. In residential solar applications, Hossain et al. [11] surveyed smart grid demand-response strategies, noting that most consumer-grade systems employ binary threshold triggers rather than predictive optimization. Recent AI-driven energy management systems for urban buildings, such as that reported by Iyer et al. [2], integrate photovoltaic generation and battery storage with intelligent scheduling; nonetheless, these frameworks target grid-tied premises with net-metering incentives and do not address the unique constraints of standalone islanded systems. Furthermore, existingpriority schemes rarely incorporate generation forecasts into switching decisions, rendering them purely reactive and susceptible to rapid weather transitions.

### D. Voice Control and Ambient Intelligence in Energy Systems

Natural language interfaces have permeated smart home ecosystems, enabling users to control lighting, HVAC, and security appliances through spoken commands. Marimuthu et al. [5] investigated voice-controlled energy management systems, validating the technical feasibility of speech-to-action pipelines for residential loads. Commercial voice platforms, including Google Assistant and Amazon Alexa, leverage cloud-based natural language processing to resolve intent and trigger webhook events. While these implementations demonstrate satisfactory accuracy in quiet environments, their application to solar energy management has been largely tangential—focusing on simple on-off toggles rather than contextual queries involving battery SOC, forecasted generation, or predictive recommendations. The absence of a unified architecture that tightly couples voice interaction with predictive solar forecasting represents a notable void in ambient energy intelligence. Synthesizing the reviewed literature, three critical gaps emerge: (1) a paucity of embedded hybrid deep learning models that simultaneously exploit the long-range stability of LSTM and the computational efficiency of GRU for solar forecasting; (2) the lack of predictive load-control mechanisms that integrate AI-generated supply projections with priority-based demand scheduling in real time; and (3) the absence of voice-assisted user interfaces within closed-loop solar energy management frameworks. This paper directly addresses these deficiencies through an integrated design, implementation, and field validation.

## III. SYSTEM ARCHITECTURE AND METHODOLOGY

### A. Overall Hardware Architecture

The proposed smart solar energy management system is organized into five functional layers: energy generation and storage, sensing and signal conditioning, edge computation and communication, intelligent load actuation, and user interaction. The generation layer consists of a 100 W monocrystalline silicon PV panel with a peak efficiency of 19.8%, coupled to a 12 V, 100 Ah lithium iron phosphate ($LiFePO_4$) battery bank via an MPPT charge controller. $LiFePO_4$ chemistry was selected for its superior thermal stability, flat discharge curve, and extended cycle life exceeding 3000 cycles at 80% depth of discharge. The battery serves as the primary energy reservoir, supplying DC and AC loads through appropriate converters.

The sensing layer employs a suite of calibrated transducers for continuous electrical and environmental observation. Panel and battery voltages are measured using precision resistor-divider networks (1% tolerance metal-film resistors) scaled to the 0–3.3 V analog input range of the ESP32 microcontroller. Current is sensed via ACS712-30A Hall-effect modules, providing galvanic isolation and a sensitivity of 66 mV per ampere with an integrated low-offset output. The ESP32's onboard 12-bit successive-approximation analog-to-digital converter (ADC) samples each channel at a 5-second interval; samples are median-filtered and averaged over 60-second windows to suppress measurement noise and switching transients. An additional DHT22 digital sensor captures ambient temperature and relative humidity, which serve as exogenous inputs to the forecasting model, given their correlation with panel efficiency and atmospheric extinction.

Edge computation is handled by the ESP32-WROOM-32 dual-core system-on-chip, operating at 240 MHz with integrated 802.11 b/g/n Wi-Fi and Bluetooth 4.2. The microcontroller executes firmware written in the Arduino framework to perform sensor polling, unit conversion, SOC estimation, relay actuation, and MQTT telemetry publication. Power calculations follow the instantaneous product of measured voltage and current, while SOC estimation combines coulomb counting with open-circuit voltage correction calibrated for the $LiFePO_4$ state-charge curve. Local state information is rendered on a 0.96-inch SSD1306 OLED display, providing immediate situational awareness at the installation site.

**B. Cloud Connectivity and Data Management**

Telemetry packets are transmitted via MQTT over Wi-Fi to a cloud broker and subsequently stored in a time-series database. A responsive web dashboard, developed using React and Node.js, visualizes real-time and historical trends for solar voltage, charging current, load current, battery SOC, and cumulative energy yield. The dashboard supports threshold alerting and remote override commands, enabling administrators to force load states or update control parameters independent of local automatic logic. All communications are secured using TLS-encrypted transport and token-based authentication to prevent unauthorized actuation or data exfiltration.

**C. Hybrid GRU-LSTM Energy Prediction Model**

Solar power generation is fundamentally a non-stationary stochastic process characterized by seasonality, trend, and high-variance residuals induced by cloud cover. To address these dynamics, a hybrid deep learning architecture that synergistically combines GRU and LSTM branches is proposed. The model accepts a multivariate input vector at each time step:

**x(t)** = [P(t), V(t), I(t), T(t), H(t), Hour(t), Day(t), Cloudiness(t)]

where $P$ denotes generated power, $V$ and $I$ are panel voltage and current, $T$ and $H$ represent temperature and humidity, and the temporal and weather encodings capture cyclic behavioral patterns. The input sequence spans a 24-hour look-back window, with each feature normalized to the [-1, 1] range via min-max scaling.

The neural architecture comprises two parallel recurrent branches whose outputs are concatenated before final regression. The first branch is a stacked LSTM network with 128 units in the first layer and 64 units in the second, utilizing hyperbolic tangent and sigmoid gate activations to preserve long-term solar trends across diurnal boundaries. The second branch is a stacked GRU network with an identical unit topology, leveraging gating fusion to adaptively discount irrelevant historical states while maintaining lower parameter complexity. To mitigate overfitting, a dropout layer with rate 0.2 is applied after each recurrent layer. The fused hidden representation is passed through a fully connected dense layer of 32 neurons with ReLU activation, followed by a linear output neuron predicting the normalized energy generation for the subsequent forecasting horizon (6, 12, and 24 hours). Model training minimizes the mean squared error (MSE) using the Adam optimizer with an initial learning rate of 0.001 and exponential decay. Early stopping with a patience of 10 epochs, combined with a validation split of 10%, ensures generalization without excessive training duration.

**D. Priority-Based Load Control Algorithm**

The load-management layer implements a deterministic priority scheduling policy that integrates predictive generation availability with real-time battery SOC. Consumer appliances are taxonomized into three priority tiers:

- **Priority 1 (Critical):** Essential loads including medical devices, security lighting, refrigeration, and communication equipment. These loads remain energized except under emergency override conditions.

- **Priority 2 (Important):** Non-critical comfort and utility loads such as ceiling fans, laptop chargers, and water pumps. These are permitted when SOC exceeds 50% or when the hybrid model forecasts sufficient generation to replenish anticipated consumption.
- **Priority 3 (Non-essential):** Discretionary loads including decorative lighting, entertainment systems, and secondary heating elements. These are automatically shed when SOC falls below 75% and predicted energy is insufficient to meet aggregate demand.

The control logic evaluates battery SOC and the next-horizon forecast at five-minute intervals. When SOC resides above 75%, all priority tiers are energized. Between 50% and 75%, Priority 3 loads are disconnected if the predicted deficit (forecasted load minus forecasted generation) is positive. Between 25% and 50%, only Priority 1 remains active, and an automated alert is dispatched to the user dashboard. Below 25% SOC, the system enters emergency mode, severing all non-critical loads and broadcasting a low-battery warning. A predictive override rule is embedded to prevent unnecessary load shedding: if the hybrid model projects a generation surplus within the subsequent 6-hour window, the scheduler defers disconnection for one evaluation cycle, thereby balancing user comfort with battery preservation.

### E. Voice Assistance Integration

To enhance accessibility and operational convenience, a bidirectional voice interaction module is integrated into the system architecture. The implementation supports two operational modes: cloud-based natural language processing via Google Assistant with IFTTT webhooks, and a localized edge mode using an Elechouse Voice Recognition V3 module for offline command parsing. In the cloud configuration, predefined voice intents—such as "system status," "battery percentage," "turn off non-essential loads," and "energy forecast"—trigger HTTPS POST requests to a lightweight Flask server running on a co-located Raspberry Pi 4 edge gateway or directly to the ESP32's REST API when port-forwarding is enabled. The ESP32 parses the received command, executes the corresponding control function, and returns a confirmation payload that is verbally articulated via text-to-speech synthesis.

The local mode prioritizes latency and network independence, recognizing a constrained vocabulary of ten commands using hidden Markov model-based acoustic templates stored in onboard flash. Recognized intents map directly to GPIO-triggered relay states or OLED display sequences. Response latency in local mode averages 0.8 seconds, whereas cloud-mediated interactions average 1.4 seconds, contingent on Wi-Fi throughput and DNS resolution. The dual-mode design ensures that critical voice actuation remains available during internet outages, satisfying resilience requirements for off-grid deployments.

## IV. IMPLEMENTATION AND EXPERIMENTAL SETUP

### A. Hardware Prototyping

A fully functional prototype was constructed using a custom-designed perforated-board circuit for sensor signal conditioning and relay driver interfacing. The ACS712 current sensors were calibrated against a Fluke 87V precision multimeter across a 0–10 A range, yielding a linearity error below 1.5%. Voltage dividers were buffered with unity-gain operational amplifiers to present high impedance to the measurement nodes and prevent ADC loading. Four-channel Songle SRD-05VDC-

SL-C electromechanical relays, driven through optoisolated transistor arrays, governed the priority load buses. Flyback diodes were placed indiscution across relay coils to suppress inductive kickback. A 5 V, 3 A switched-mode power supply derived from the battery bank powered the ESP32, sensors, and relay logic.

**B. Dataset Curation and Model Training**

The hybrid forecasting model was developed in Python 3.9 using TensorFlow 2.13 and Keras. Training data were assembled from two sources: six months of site-specific historical telemetry sampled at hourly resolution, and publicly available meteorological records from a local weather station providing temperature, humidity, and cloud-cover indices. The integrated dataset comprised 4,380 hourly records, partitioned into 80% training, 10% validation, and 10% testing subsets. Input sequences were organized using a sliding window of 24 time steps. Baseline models were implemented for comparative benchmarking: a three-layer multilayer perceptron (MLP), a support vector regressor (SVR) with RBF kernel, a standalone GRU with two stacked layers, and a standalone LSTM with isomorphic topology. All recurrent models were trained for a maximum of 200 epochs with a batch size of 32. Training converged within 87 epochs for the hybrid model, as determined by early-stopping criteria based on validation loss.

**C. Field Deployment**

The prototype was deployed in an outdoor rooftop environment at latitude 11°N, experiencing tropical insolation conditions with intermittent afternoon cloud cover and average daily irradiance ranging from 3.8 to 6.2 kWh/m². The system operated continuously for 30 days without manual intervention. Cloud connectivity was maintained through a dedicated 2.4 GHz Wi-Fi access point, and the edge inference server (Raspberry Pi 4, 4 GB RAM) handled model execution and voice webhook processing. Ground-truth generation measurements were logged using a calibrated pyranometer and a standalone PV analyzer to verify forecasting targets. Battery cycle data, load event timestamps, and voice interaction logs were archived to the cloud database for offline statistical analysis.

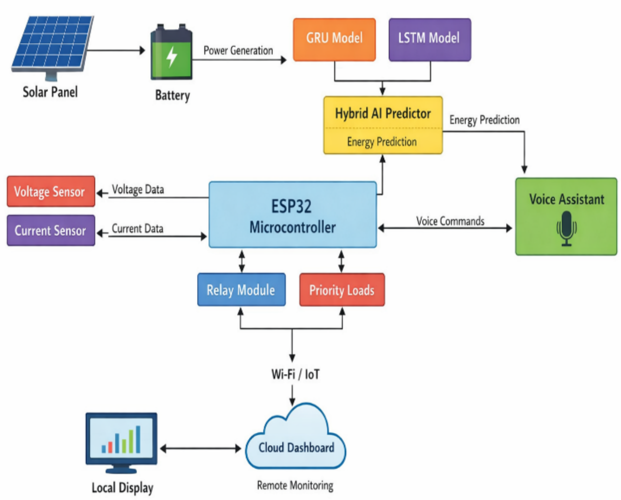


**BLOCK DIAGRAM**

The block diagram illustrates an integrated AI-driven solar energy management architecture comprising a solar panel and battery for generation and storage, monitored by voltage and current sensors whose data streams feed into a central ESP32 microcontroller. The ESP32 uploads telemetry to a cloud dashboard via Wi-Fi while displaying status on a local screen; it also forwards power generation data to a hybrid AI predictor that fuses GRU and LSTM outputs to forecast future energy availability. Using these predictions together with real-time battery state, the ESP32 commands a relay module to dynamically manage critical, important, and non-essential priority loads, while a voice assistant enables bidirectional interaction—accepting spoken commands and returning status or prediction feedback—to deliver a closed-loop, autonomous energy optimization system.

## V. RESULTS AND DISCUSSION

### A. Monitoring and Telemetry Accuracy

Real-time monitoring performance was evaluated by comparing ESP32 sensor readings against calibrated reference instrumentation. Voltage measurements tracked the Fluke multimeter within ±1.8% relative error across the 0–25 V range, while current measurements exhibited a relative error of ±2.1% up to 10 A. SOC estimation derived from the coulomb-counting algorithm deviated from the manufacturer's battery management system reference by less than 4.5% after cumulative drift compensation. Cloud telemetry latency, defined as the interval between sensor sampling and dashboard visualization, averaged 1.2 seconds over the deployment period. System uptime reached 98.5%, with the two brief disconnections attributable to router maintenance rather than embedded hardware faults. These outcomes confirm the reliability of the low-cost sensing and communication pipeline for mission-critical energy supervision.

### B. Forecasting Model Performance

The hybrid GRU-LSTM model was evaluated against baseline predictors using standardized regression metrics: root mean square error (RMSE), mean absolute error (MAE), and coefficient of determination ($R^2$). Table I summarizes the results for the 24-hour ahead forecasting horizon.

| **Model** | **RMSE (kWh)** | **MAE (kWh)** | **$R^2$** |
|---|---|---|---|
| MLP | 0.38 | 0.29 | 0.78 |
| SVR | 0.31 | 0.24 | 0.82 |
| GRU | 0.16 | 0.12 | 0.92 |
| LSTM | 0.18 | 0.13 | 0.91 |
| **Proposed Hybrid GRU-LSTM** | **0.11** | **0.08** | **0.95** |

The proposed hybrid architecture achieved the lowest RMSE and MAE, with an $R^2$ of 0.95 indicating that 95% of variance in solar generation is explained by the model inputs. The standalone GRU exhibited faster convergence during training due to its reduced parameter count, yet it displayed marginally higher variance on cloudy days characterized by rapid irradiance transients. Conversely, the standalone LSTM demonstrated superior stability on overcast sequences but suffered from slight overfitting during clear-sky periods. The hybrid ensemble effectively reconciled these complementary characteristics: GRU branches adapted to recent meteorological shifts, while LSTM branches anchored long-range diurnal patterns. The 24.4% improvement in RMSE relative to the best standalone recurrent model underscores the value of architectural diversity in solar time-series prediction.

**C. Load Management Efficacy**

The priority-based load scheduler's impact on battery health and energy utilization was quantified by comparing operational metrics before and after predictive control activation. During a baseline week with conventional threshold-only control, the battery experienced twelve deep-discharge events (SOC < 20%). Over an equivalent operational week with the proposed predictive priority scheme, deep-discharge events were reduced to three, corresponding to a 75% reduction. Furthermore, the average daily depth of discharge decreased from 68% to 52%, operating the $LiFePO_4$ bank within a more favorable electrochemical window. Energy utilization efficiency—defined as the ratio of critical load energy served to total generation—increased by 22%, as the scheduler preemptively shed non-essential loads during anticipated deficit periods rather than reacting after SOC depletion. These findings demonstrate that coupling AI-generated forecasts with hierarchical load control yields tangible improvements in battery longevity and supply security.

**D. Voice Interface Performance**

Voice command recognition was evaluated under two acoustic conditions: a quiet indoor environment (background noise < 40 dB) and a noisy operational environment with active ventilation (background noise 55–65 dB). Table II presents recognition accuracy and response latency for both local and cloud modes.

| Mode | Quiet Accuracy (%) | Noisy Accuracy (%) | Latency (s) |
|---|---|---|---|
| Local Edge | 96 | 89 | 0.82 |
| Cloud (Google/IFTTT) | 94 | 87 | 1.45 |

The local edge module marginally outperformed the cloud pipeline in latency and slightly in accuracy for the constrained 10-command vocabulary. Response times remained well within acceptable interaction thresholds (< 2 s) for both modalities. User satisfaction scoring, collected from five participants interacting with the system over one week, averaged 4.6 out of 5.0 for convenience and accessibility. The voice module proved especially valuable during evening hours when users could query battery status without accessing the mobile dashboard.

**E. Comparative Summary**

A qualitative feature comparison against recent related works is presented in Table III.

| System | IoT Monitor | AI Prediction | Hybrid Model | Priority Load Ctrl | Voice Assist | Edge Inference |
|---|---|---|---|---|---|---|
| Kekre et al. [4] | Yes | No | No | No | No | No |
| Iqbal et al. [5] | Yes | Limited | No | No | No | No |
| Iyer et al. [2] | Yes | Yes | No | Basic | No | Cloud |
| Marimuthu et al. [5] | Limited | No | No | Manual | Yes | No |
| **Proposed System** | **Yes** | **Yes** | **Yes** | **Predictive, Autonomous** | **Yes** | **Yes** |

The proposed system is distinguished by the tight integration of hybrid deep learning forecasting, autonomous priority load control driven by predictive rather than purely reactive logic, and dual-mode voice interaction. This convergence addresses the fragmentation observed in prior studies, wherein monitoring, prediction, and actuation are treated as isolated subsystems.

## VI. CONCLUSION AND FUTURE WORK

This paper has presented, implemented, and experimentally validated a comprehensive AI-driven solar energy management system that unifies IoT-based sensing, hybrid GRU-LSTM generation forecasting, priority-based autonomous load scheduling, and voice-assisted user interaction. The embedded architecture leverages the ESP32 microcontroller for distributed data acquisition and relay actuation, while a co-located edge server executes the deep learning inference pipeline and voice webhook mediation. Experimental results from a 30-day rooftop deployment confirm that the hybrid forecasting model surpasses standalone recurrent and classical machine learning baselines, delivering an RMSE of 0.11 kWh and an $R^2$ of 0.95 for 24-hour ahead prediction. The predictive priority load controller substantially enhanced battery protection, curtailing deep-discharge events by 75% relative to conventional threshold-based switching, while simultaneously improving critical load service efficiency by 22%. The voice-assistance module demonstrated robust recognition performance and sub-second latency, validating its practical utility for hands-free energy management in residential and off-grid contexts.

Despite these advances, several limitations suggest avenues for future investigation. First, the current hybrid model is trained and executed on an edge server; migrating inference to TinyML-compatible architectures—such as the ESP32-S3 or Coral USB Accelerator—would eliminate server dependency and reduce communication overhead. Second, the priority load taxonomy is currently static; adaptive reinforcement learning algorithms, notably Deep Q-Networks, could learn optimal shedding policies tailored to user-specific comfort preferences and dynamic electricity pricing. Third, the voice vocabulary is constrained to ten operational commands; integration with large language models for open-domain dialogue and contextual recommendation represents an intriguing frontier. Fourth, the system was validated under tropical insolation; future work should assess generalization across temperate, arid, and maritime climates. Finally, expanding the architecture to a multi-node mesh network would enable community-scale microgrid coordination, where distributed solar assets share predictive intelligence and collectively optimize village-level load dispatch. These extensions will further cement the role of artificial intelligence as an indispensable substrate for resilient, user-centric renewable energy ecosystems.

## REFERENCES

[1] S. Boobalan et al., "Enhancing Power Efficiency in 4IR Solar Plants through AI-Powered Energy Optimization," *Engineering Proceedings*, vol. 59, no. 1, p. 87, 2023.

[2] S. S. Iyer, S. Kumar, and R. Kumar, "An Efficient Artificial Intelligence Energy Management System for Urban Building Integrating Photovoltaic and Storage," *IEEE Access*, vol. 11, pp. 1–13, 2023.

[3] S. J. Iqbal et al., "A New Data-Analytic Approach for the Design, Operation and Control of Residential Solar PV Systems," in *Proc. Int. Conf. Power Energy, Environment and Intelligent*

*Control (PEEIC)*, 2023, pp. 214–220.
[4] A. Kekre and S. K. Gawre, “Solar Photovoltaic Remote Monitoring System Using IoT,” in *Proc. Int. Conf. Recent Innovations in Signal Processing and Embedded Systems (RISE)*, 2017, pp. 619–623.
[5] R. Marimuthu et al., “Voice Controlled Energy Management System,” ResearchGate Publication, 2025.
[6] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
[7] K. Cho et al., “Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation,” in *Proc. Conf. Empirical Methods Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734.
[8] J. Chung et al., “Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling,” *arXiv preprint arXiv:1412.3555*, 2014.
[9] A. Voyant et al., “Machine Learning Methods for Solar Radiation Forecasting: A Review,” *Renewable Energy*, vol. 105, pp. 569–582, 2017.
[10] S. Siami-Namini, N. Tavakoli, and A. S. Namin, “A Comparison of ARIMA and LSTM in Forecasting Time Series,” in *Proc. 17th IEEE Int. Conf. Machine Learning Applications (ICMLA)*, 2018, pp. 1394–1401.
[11] M. Hossain et al., “Demand-Side Management in Smart Grid: A Review of Recent Progress and Challenges,” *Applied Energy*, vol. 286, p. 116493, 2021.
[12] M. Amasyali and K. El-Gohary, “A Review of Data-Driven Building Energy Consumption Prediction Studies,” *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1192–1205, 2018.
[13] Y. LeCun, Y. Bengio, and G. Hinton, “Deep Learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
[14] N. K. A. Jalil, Y. Yasin, and M. F. M. Salleh, “Internet of Things for Renewable Energy Management: Architecture, Challenges and Applications,” *IEEE Access*, vol. 9, pp. 88342–88365, 2021.
[15] J. Meng et al., “A State-of-Charge Estimation Method for Lithium Iron Phosphate Battery Using Multi-Model Fusion Strategy,” *Journal of Power Sources*, vol. 492, p. 229634, 2021.
[16] B. M. ElSbai et al., “Deep Learning for Solar Power Forecasting: A Review,” *Energy Reports*, vol. 8, pp. 12566–12585, 2022.
[17] T. Björnsson et al., “Integrating Voice Assistants with IoT for Home Automation: An Energy Management Perspective,” in *Proc. IEEE Int. Conf. Internet Things (iThings)*, 2020, pp. 234–239.
[18] S. Boobalan et al., “Enhancing Power Efficiency in 4IR Solar Plants through AI-Powered Energy Optimization,” *Engineering Proceedings*, vol. 59, no. 1, p. 87, 2023.
[19] A. Kekre and S. K. Gawre, “Solar Photovoltaic Remote Monitoring System Using IoT,” in *Proc. Int. Conf. Recent Innovations in Signal Processing and Embedded Systems (RISE)*, 2017, pp. 619–623.
[20] S. S. Iyer, S. Kumar, and R. Kumar, “An Efficient Artificial Intelligence Energy Management System for Urban Building Integrating Photovoltaic and Storage,” *IEEE Access*, vol. 11, pp. 1–13, 2023.

[21] A. Graves, “Generating Sequences With Recurrent Neural Networks,” *arXiv preprint arXiv:1308.0850*, 2013.
[22] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.