

AN AI-AUGMENTED HYBRID RENEWABLE ENERGY SYSTEM FOR INDIAN RAILWAY STATIONS: A CASE STUDY AT MUMBAI CENTRAL TERMINUS

MOHAMMED AFWAN

MASTER OF ENGINEERING IN RENEWABLE ENERGY

Abstract - The Indian Railways consumes about 20 billion units of electricity in a year, constituting about 2.4 percent of the national electricity demand. The Railway Ministry has pledged to achieve net-zero carbon emissions from its operations by 2030. Most of the electricity consumed by Railways is for traction. However, considerable demand exists at the station level. Therefore, in the absence of any traction-related savings, an on-site renewable energy system at major railway stations can help the Ministry achieve its net-zero carbon emissions goal. The present study focuses on sizing and simulating an integrated renewable energy system at Mumbai Central Terminus. The system is designed to incorporate 1.41 MWp of rooftop solar PV, piezoelectric energy harvesting from footfalls of passengers, piezoelectric energy harvesting from rail vibrations, and thermoelectric energy recovery from the waste heat of the air conditioning system. The system is buffered by a 3.1 megawatt-hour lithium-iron-phosphate battery and connected through a 1.5 megavolt-ampere inverter to the grid. Each subsystem is sized from first principles using internationally accepted methods, with all governing equations and substituted values reported explicitly so the calculations can be reproduced. The same physical hardware is then simulated under two control strategies in MATLAB and Simulink. The conventional version uses today's standard controllers, while an artificial-intelligence-augmented version replaces them with a long short-term memory forecaster, neural-network maximum-power-point trackers, a Twin-Delayed Deep Deterministic Policy Gradient reinforcement-learning energy manager, and an XGBoost-with-One-Class-Support-Vector-Machine predictive maintenance pipeline. Across a 24-hour and 30-day benchmarking horizon, the AI layer raised annual renewable yield by 5.9 percent, reduced grid imports by 26 percent, cut daily battery cycling by 23 percent, and provided multi-week warning of four common failure modes. Renewable share of station load grew from 51 percent to 58 percent, all at zero additional capital cost on the physical plant. Replicated across more than two thousand already solarised stations on the Indian Railways network, the same software-only upgrade represents roughly 270 gigawatt-hours of additional renewable energy each year, comparable to a 150-megawatt utility-scale solar plant, with no extra panels, land, or batteries.

Keywords: *renewable energy; railway station; hybrid microgrid; reinforcement learning; predictive maintenance*

INTRODUCTION

India's railway network represents the largest electricity consumer within the country. Indian Railways carries roughly 8 billion passenger journeys a year, runs over 13,000 trains daily, and operates more than 7,300 stations. The annual electricity consumption of Indian Railways accounts for approximately 2.4 percent of the nation's total electricity demand, exceeding 20 billion kWh per year [1]. While most of this energy is allocated to traction, a substantial proportion is consumed within station premises for lighting, escalators, ticketing systems, public-address systems, signalling, and air conditioning. Lighting, escalators, ticketing systems, public-address systems, signalling, and (increasingly) air conditioning all run continuously, and a medium-to-large terminal can pull anywhere between 100,000 and 500,000 kWh in a single month. [1].

Almost all of that electricity comes from the national grid. The Central Electricity Authority's most recent baseline database, Version 21, published in November 2025, places the all-India weighted-average grid emission factor at 0.7117 kg of CO₂ per kWh delivered [2]. A 500,000 kWh-per-month station is therefore directly responsible for over 4,200 tonnes of CO₂ a year, before transmission and distribution losses are added.[2]

In this context, the Ministry of Railways has pledged to declare its tracks as a "net zero carbon" by 2030 [3]. Two complementary approaches are needed to achieve this target: increasing renewable generation capacity to match the network demand and improving the efficiency of the renewable systems already in place. This is the latter, which is addressed in the present study, where the additional useful energy that can be obtained by software-based optimization [3] is quantified.

The demand at the station-level cannot be met by a single renewable energy source. Solar PVs can only be used during the day and cannot produce energy when energy demand is high at night. Piezoelectric systems are able to continuously generate energy from the passage of pedestrians, yet only in small amounts. Waste heat from HVAC condensers can be utilized by thermoelectric generators (TEGs) for generating electricity, but this is only possible during the operating time of the HVAC cooling systems. This is why these sources have to be coupled to energy storage and intelligent control systems. When the three of them are combined with a battery to offset supply-and-demand energy gaps, and a smart controller to flexibly manage them, the total energy capacity greatly exceeds the demand that the grid could supply.

This research presents the design of a hybrid system at Mumbai Central Terminus and, more specifically, examines the question of how much additional useful energy can be obtained from existing physical system components through a software upgrade. To answer this question, two control variants are built. The first control variant is built using today's standard practices. The second control variant is built using an AI layer that includes a long short-term memory (LSTM) forecasting model, a neural network-based MPPT, a Twin-Delayed Deep Deterministic Policy Gradient (TD3) reinforcement-learning-based battery dispatch agent, and an XGBoost-plus-One-Class-SVM predictive maintenance pipeline. All of these run in MATLAB and Simulink using the same set of inputs.

This study is exclusively based on simulations, as no physical prototype has been created. The target isn't a design ready for deployment, but rather a retrospective study to understand the engineering possibilities of incorporating an AI layer into railway stations with existing physical infrastructure. Currently, the Indian Railways has more than 2,000 operational stations. Given this large number, the rationale for this type of incremental, software-only optimization is justified.

ENERGY LANDSCAPE AND POLICY CONTEXT

2.1 STATION-LEVEL ENERGY DEMAND

There is a scarcity of dependable data on non-traction station consumption in India. Xu et al.'s findings are valuable in that regard. They assess HVAC consumption in the context of Chinese railway stations. They established that in terminal stations with conditioned waiting areas, air conditioning consumption can account for as much as 40% (or more) of total station energy. [4]. A separate study of an underground station in Bilbao reported that lighting consumes roughly a third of all energy used. [5]. India's climate, with peak summer temperatures regularly above 35 degrees Celsius and a wet monsoon that spikes humidity, makes both contributions large. Combining these patterns with the limited Indian load data available yields an estimate of 100,000 to 500,000 kWh per month for medium-to-large terminal stations. [4].

2.2 POLICY CONTEXT: NET-ZERO BY 2030

The Ministry of Railways committed to net-zero by 2030 in 2020, and progress has been substantial. As of November 2025, **812 MW of solar capacity** and **93 MW of wind capacity** had been commissioned for traction needs, with another 1,600 MW of round-the-clock renewable power tied up under power-purchase agreements. [3]. More than 2,000 stations now run on rooftop or trackside solar, and a handful have received the BEE's Shunya net-zero certification. The longer-term ambition is 20 GW of railway-owned solar capacity by 2030. Solar adoption at the station level has moved from pilot to scale; what has not moved at scale is integration. Most stations run a single source, most lack

significant on-site storage, and almost none use intelligent controls beyond off-the-shelf inverter logic. [6].

2.3 THE THREE SOURCES CONSIDERED

2.3.1 Solar photovoltaics

A solar cell turns sunlight directly into electricity by absorbing photons in a semiconductor and separating the resulting charge carriers across the built-in electric field at the p-n junction [6]. Modern crystalline-silicon modules deliver 20-22 percent module efficiency, and well-designed tropical rooftop systems deliver around 75-85 percent of that nameplate efficiency at the AC delivery point after accounting for inverter losses, soiling, temperature, and wiring [7]. Mumbai's annual irradiance of about 2,000 kWh per square metre per year makes rooftop solar by far the dominant generation source in any plausible hybrid for this site. [7].

2.3.2 Piezoelectric harvesting

When you compress certain crystals (quartz, lead zirconate titanate, polyvinylidene fluoride), they generate a small voltage. This is the direct piezoelectric effect, first reported by the Curie brothers in 1880. At a railway station, two distinct opportunities exist: passenger footfall on platform tiles, where each footstep transfers 2-7 J of mechanical energy, of which 0.5-2 mWh can be electrically recovered in field conditions; and mechanical vibration from train passage, where cantilever PZT benders mounted between sleepers along the rails produce up to 120 kW per kilometre of instrumented track during dense traffic. Pavegen, the British company whose V3 floor tile is the dominant commercial product in this category, has installed kinetic flooring at Heathrow Terminal 3 and the Bird Street "smart street" installation in London's West End [10], [43]. Both yields are modest in absolute terms, but they keep producing through the night and through the monsoon, when solar collapses. [9].[8][9] [10][11][12]

2.3.3 Thermoelectric generation

A thermoelectric generator uses the Seebeck effect. [13]. Hold one side of a semiconductor pair hotter than the other, and a small voltage develops across it. A functional

TEG module positions many of these junctions thermally in parallel and electrically in series. Regarding HVAC systems, the waste heat expelled by the condenser is a continuous and predictable source. The strength of the waste heat correlates directly with the cooling demand. For the cooling demand in the HVAC system, each 1 kW of mechanical input to the compressor generates approximately 1.3 kW of waste heat. Bismuth telluride (Bi₂Te₃) modules, the workhorse of near-ambient thermoelectric work, achieve 1-2 percent conversion efficiency at the typical 20-30 K gradients available from tropical condenser exhausts. [14]. [12].

2.4 WHY THE CONTROLLER MATTERS

Even where renewable sources are deployed at scale, the controls layer often stays unsophisticated. MPP trackers in commercial rooftop inverters mostly use Perturb-and-Observe or a close variant. Energy management between sources, where multiple sources exist, is rule-based with hard priority lists. Predictive maintenance, where it exists, is calendar-based rather than condition-based. The result is a deployed system whose realised yield typically falls 5-15 percent below the theoretical maximum its hardware could deliver under intelligent control [15]. Closing that gap is what the AI layer in this research is designed to do. [21].

SYSTEM ARCHITECTURE AND METHODOLOGY

3.1 HYBRID SYSTEM OVERVIEW

Figure 3.1 shows the architecture of the hybrid system at Mumbai Central. Four power conditioning stages tie three generation sources to a common 400 V DC bus. A 3.1 MWh lithium-iron-phosphate battery hangs off the bus through a bidirectional buck-boost converter, and a 1.5 MVA inverter delivers the result to the station's AC switchboard. The four components of AI (the LSTM forecaster, the ANN-based MPPT controllers, the TD3 reinforcement-learning energy manager, and the XGBoost with OC-SVM for predictive maintenance) "float" at the top of the system. They communicate with the physical hardware through precisely designed interfaces.

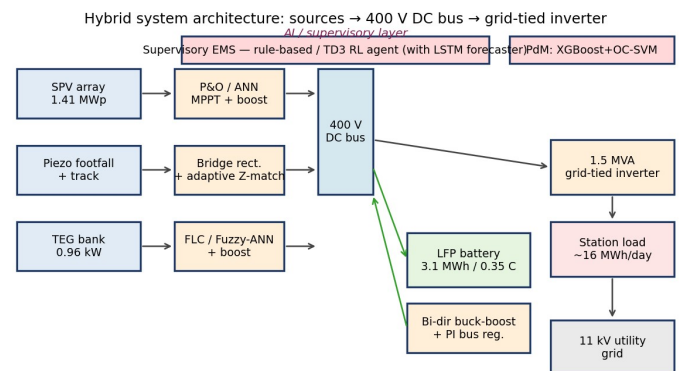


Figure 3.1 Architecture of the hybrid SPV-piezoelectric-TEG system with an AI supervisory layer.

The overlay arrangement matters. No power-electronic component, sensor, battery cell, or module is changed when the AI layer is added; only the controller running on top is. That makes the AI a clean retrofit on existing solarised stations, and it means the worst case of any AI fault is a fall-back to the rule-based baseline rather than a hardware loss.

3.2 METHODOLOGY

Two Simulink models of the same system are built. In the first, the controllers are conventional. In the second, they are AI-augmented. Both are simulated on identical site inputs (Mumbai climate from NASA POWER, Mumbai Central load profile constructed from comparable-station benchmarks, train timetable, HVAC operating schedule) and benchmarked against each other on five metrics: MPPT tracking efficiency, forecast accuracy, battery cycling, grid imports, and predictive-maintenance lead time. Sizing for every subsystem follows a first-principles approach. Each chapter that follows writes the governing equations out explicitly, plugs in the site values, and arrives at the chapter's sizing decision. Where I made a design choice (the series-parallel layout of the PV array, the depth-of-discharge limit on the battery, the heat-exchanger capture fraction for the TEG bank), I state the choice and explain the reasoning rather than simply quoting a number.

SITE, CLIMATE, AND LOAD PROFILE

4.1 REFERENCE SITE: MUMBAI CENTRAL TERMINUS

Mumbai Central Terminus (station code MMCT) is a major terminal on the Western line of Mumbai's suburban railway, and the originating terminus for several long-distance services. Its coordinates are 18.97°N, 72.82°E. According to Indian Railways' 2024 station-classification data, Mumbai Central recorded around 22 million passengers that year, equivalent to roughly 60,000 per day across both suburban

and long-distance platforms. [16]. Three site features make this station a good reference. The main platform-canopy roof spans approximately 14,150 square metres, measured from publicly available satellite imagery. The hard-tiled concourses see meaningful pedestrian throughput, useful for piezoelectric analysis. And a centrally chilled HVAC system serves the enclosed waiting halls and admin blocks, providing a continuous waste-heat source for the thermoelectric stage. [1].

4.2 CLIMATE AND SOLAR IRRADIANCE

Mumbai's climate is tropical wet-and-dry, with a pronounced south-west monsoon between June and September. Long-term meteorological data for the site were pulled from the NASA Prediction of Worldwide Energy Resources database for the $0.5^\circ \times 0.5^\circ$ grid cell centred on the station. [17] and cross-checked against the Indian Meteorological Department's 1991-2020 climatological normals [18]. Figure 4.1 plots the resulting monthly average daily Global Horizontal Irradiance. [14,15].




Figure 4.1 Mumbai monthly average daily Global Horizontal Irradiance.

The annual mean is 5.3 kWh per square metre per day. Irradiance peaks in March-May at around 7.0 kWh per square metre per day under clear post-winter skies and falls to 3.4-3.6 in July-August during the monsoon. That 35-50 percent drop for three months of the year is the defining feature of any Mumbai PV design. The optimal fixed-tilt angle for a Mumbai-latitude installation is approximately 18 degrees south, derived from the standard rule that the optimum annual fixed-tilt equals the site latitude. [19]. [16].

4.3 STATION LOAD PROFILE

A complete metered load curve for Mumbai Central is not in the public domain. The profile used in this research is constructed from three sources: published energy-consumption breakdowns for comparable large terminal stations. [4] [5], typical commercial-building load shapes for Mumbai's climate scaled to the estimated connected load, and the station's known operating schedule, including the timetable of train arrivals and departures.




Figure 4.2 Mumbai Central representative weekday: load profile and SPV generation overlay

Three regions appear across the day. Overnight (00:00-04:00), the load is dominated by platform lighting, signage, public-address systems, and idle-mode HVAC, totalling around 290-320 kW. From 06:00, the load rises sharply as morning suburban services begin and HVAC ramps up. A first peak at 980-990 kW is reached in the late morning. Between 18:00 and 19:30, a secondary peak, fairly higher than the first, between 1,080-1,100 kW, represents the evening peak hour for passengers, lighting, and the continued demand for cooling. Considering this, the total daily consumption of the station's electrical supply for cooling demand over the entire year would be estimated at about 16,200 kWh, leading to an annual estimated cooling demand of about 5.9 GWh.

SOLAR PHOTOVOLTAIC SUBSYSTEM

5.1 THE PHOTOVOLTAIC EFFECT AND SINGLE-DIODE MODEL

A solar cell is a thin slice of semiconductor, almost always silicon, with a built-in electric field at the junction where the doping changes. When a photon is absorbed, it kicks an electron loose, creating a mobile electron and a positively charged hole. The built-in field separates the two before they recombine, pushing the electron toward one terminal and the hole toward the other. Connect a load between the terminals, and a current flows. The effect was first reported by Becquerel in 1839, and the device physics is treated comprehensively in Würfel and Würfel. [6]. [7].

The terminal current-voltage relationship of a real cell is described by the single-diode equivalent-circuit model. Figure 5.1 shows the equivalent circuit, and Equation 5.1 shows the governing relation:




Figure 5.1 Single-diode equivalent circuit of a solar cell.

$$I = I_{ph} - I_0 \left[\exp\left(\frac{V + I \cdot R_s}{n \cdot V_T}\right) - 1 \right] - (V + I \cdot R_s) / R_{sh} \quad (5.1)$$

where I_{ph} is the light-generated current (proportional to irradiance), I_0 is the diode reverse-saturation current, n is the diode ideality factor (typically 1.0-1.5 for crystalline silicon), V_T is the thermal voltage equal to $k \cdot T / q$ (about 0.026 V at 25 °C), R_s is the series resistance, and R_{sh} is the shunt resistance. Three operating points define the behaviour of any module: the short-circuit current I_{sc} (where $V = 0$), the open-circuit voltage V_{oc} (where $I = 0$), and the maximum power point at coordinates (V_{mp}, I_{mp}) where the product $V \cdot I$ is maximised. Figure 5.2 shows the I-V and P-V characteristics of the reference module at four irradiance levels.

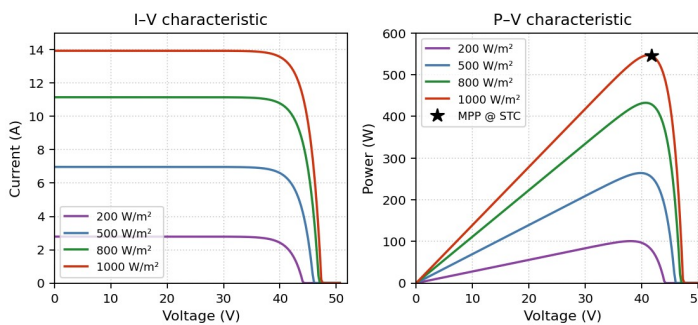


Figure 5.2 I-V and P-V characteristics of the reference 545 Wp monocrystalline module at four irradiance levels (cell temperature 25 °C).

5.2 MODULE TECHNOLOGY SELECTION

I evaluated three types of commercial photovoltaic technologies: mono- and multi-crystalline silicon and thin-film (copper indium gallium selenide, and Cadmium Telluride) technologies. This site has three specific reasons for choosing mono-crystalline modules. The available roof area is limited, so higher module efficiency equates to more installed capacity in the same footprint. The temperature coefficient for mono-crystalline silicon, although not the best of the three, is adequate at normal Mumbai cell temperatures with some form of rear ventilation. And monocrystalline has fallen below

polycrystalline on a rupees-per-Wp basis in the Indian market since 2021, removing the historical economic argument for polycrystalline. The reference module adopted is an Adani Solar Elan Series ASM-M10-144-545 module (545 Wp, monocrystalline PERC, Indian-manufactured, MNRE ALMM-listed) [41]. Its standard test condition parameters are $V_{oc} = 49.65$ V, $I_{sc} = 13.94$ A, $V_{mp} = 41.94$ V, $I_{mp} = 13.01$ A, module efficiency 21.14 percent, and temperature coefficient of P_{max} equal to -0.34 percent per °C. Module dimensions are 2.279 m x 1.131 m, giving a module area of 2.578 square metres. [17],[20] [21]

5.3 ARRAY SIZING USING THE IEC 61724-1 METHOD

The IEC 61724-1 Performance Ratio method [7] is the international standard for sizing rooftop and utility-scale PV systems. The annual energy yield is given by:

$$E_{ANNUAL} = P_{INSTALLED} \times H_{POA} \times PR \quad (5.2)$$

where $P_{INSTALLED}$ is the installed DC capacity (kWp), H_{POA} is the annual plane-of-array irradiance (kWh per square metre per year), and PR is the dimensionless Performance Ratio that aggregates every system loss between the module nameplate and the AC delivery point. Each input is calculated from a chain of physical measurements and standard correction factors.

5.3.1 Available array area

The roof area, measured from satellite imagery, is 14,150 square metres. NREL's O&M Best Practices guidance [22] recommends reserving 25 percent of the roof area for walkways, equipment plinths, and access easements:

$$A_{USABLE} = A_{ROOF} \times (1 - f_{ACCESS}) = 14,150 \times (1 - 0.25) = 10,612.5 \text{ m}^2 \quad (5.3)$$

Industrial flat-roof Ground Coverage Ratios sit between 0.65 and 0.75 to allow inter-row spacing for shading. Using 0.70:

$$A_{ARRAY} = A_{USABLE} \times GCR = 10,612.5 \times 0.70 = 7,428.75 \text{ m}^2 \quad (5.4)$$

5.3.2 Installed capacity

For modern monocrystalline modules at 21 percent efficiency on industrial-roof racking, the array-level power density is approximately 0.190 kWp per square metre. [20]:

$$P_{\text{INSTALLED}} = A_{\text{ARRAY}} \times \rho_{\text{KWp}} = 7,428.75 \times 0.190 = 1,411.5 \text{ kWp} \quad (5.5)$$

The required module count is:

$$N_{\text{MODULES}} = \lceil P_{\text{INSTALLED}} \times 1000 / P_{\text{MODULE}} \rceil = \lceil 1,411,500 / 545 \rceil = 2,590 \text{ modules} \quad (5.6)$$

5.3.3 Series-parallel layout — design choice

With 2,590 modules and a 400 V DC bus, the array can be arranged as N_s modules in series per string and N_p parallel strings, subject to the constraint:

$$N_s \times N_p = N_{\text{MODULES}} = 2,590 \quad (5.7)$$

The series count is chosen so that the string voltage at the maximum power point sits at or just above the DC bus voltage, allowing the boost-stage MPPT to pull it down to bus level. The string voltage at MPP is $N_s \times V_{\text{mp}} = N_s \times 41.85 \text{ V}$. For the 400 V bus, the smallest N_s that places the string above the bus is:

$$N_s = \lceil V_{\text{BUS}} / V_{\text{mp}} \rceil = \lceil 400 / 41.94 \rceil = 10 \text{ modules in series} \quad (5.8)$$

With $N_s = 10$, the string Voc is $10 \times 49.65 = 496.5 \text{ V}$, comfortably within the 600 V SELV limit common to commercial inverters, and the string Vmp is $10 \times 41.94 = 419.4 \text{ V}$, sitting just above the 400 V bus. The number of parallel strings follows from the constraint:

$$N_p = N_{\text{MODULES}} / N_s = 2,590 / 10 = 259 \text{ parallel strings} \quad (5.9)$$

The specific arrangement of 10 series \times 259 parallel is intentional. If I used a higher series count, say 12 or 15 modules, the string Vmp would exceed 500 V, and a greater derating margin on the inverter input would be required. If I used fewer series counts, say 5 or 7 modules, the boost stage would have to ramp up a small voltage by a large ratio. This

would be expensive and would require a copper wire with a large diameter to handle the higher per-string current. Ten in series is the optimal arrangement, as the string Vmp is close to bus voltage, parallel current per string is able to maintain an acceptable Imp of 13.03 A, and no modules are left unassigned when the array is divided into 259 strings. The total array short-circuit current at standard test conditions is $N_p \times I_{\text{sc}} = 259 \times 13.94 = 3,610.5 \text{ A}$, and the array MPP current is $259 \times 13.01 = 3,369.6 \text{ A}$.

5.3.4 Plane-of-array irradiance

Annual horizontal irradiance for Mumbai is $H_{\text{GHI}} = 5.3 \text{ kWh per square metre per day}$, integrating to:

$$H_{\text{GHI,ANNUAL}} = H_{\text{GHI}} \times 365 = 5.3 \times 365 = 1,934.5 \text{ kWh/m}^2/\text{yr} \quad (5.10)$$

For a south-facing 18° tilt at 19°N, the plane-of-array irradiance is approximately 4 percent higher than the horizontal value. [19]:

$$H_{\text{POA}} = H_{\text{GHI,ANNUAL}} \times f_{\text{TILT}} = 1,934.5 \times 1.04 = 2,012 \text{ kWh/m}^2/\text{yr} \quad (5.11)$$

5.3.5 Performance Ratio

PR is the product of independent loss factors:

A note on the PR formulation used here. In IEC 61724-1 [8], there are two formulations of PR, which are defined operationally. The first version, the measured PR, is the ratio of the delivered AC yield to the reference yield, $P_{\text{STC}} \cdot H_{\text{POA}} / G_{\text{STC}}$. This version combines all the losses occurring in the system into a single numerical value. The second version, the predicted PR, is defined by the product of independent loss mechanisms, as shown in Table 1. The two formulations have important differences. The expressed cascade of loss mechanisms in predicted PR can under- or over-estimate the measured PR, due to the overlap of individual mechanisms (for example, module mismatch, soiling, etc., can be defined at different levels and can also be counted, in part, cumulatively). The 0.789 value used in this research is the cascade version of the predicted PR and corresponds to the range of 0.78-0.82 found at similar

tropical-rooftop installations [20]. It is assumed that the measured PR of the deployed system will be in that range. Should this not be the case, the predictive maintenance pipeline in Chapter 9 will explore individual factors that could account for the discrepancy.

$$PR = \eta_{inv} \times \prod (1 - L_i) \quad (5.12)$$

Table 1: Performance Ratio loss components for the Mumbai Central rooftop installation.

Component	Symbol	Value	Source
Inverter efficiency	η_{inv}	0.970	Modern string inverters
Temperature loss	L_{TEMP}	10%	Skoplaki and Palyvos [18] [23]
Soiling loss	L_{SOIL}	3%	Kazem et al. [19] [24]
DC + AC wiring loss	L_{WIRE}	2%	Industry standard
Module mismatch+ LID	$L_{MISMATCH}$	2%	IEC 61724-1 [8] [7]
Near-shading loss	L_{SHADE}	2%	Site-specific
Availability loss	L_{AVAIL}	1%	Industry typical

A 10 percent loss in temperature in Table 1 results in an estimated average cell temperature of about 29 K above STC at the module temperature coefficient of -0.34 percent/K. This is likely the highest in the range for a Mumbai flat-roof installation when using the assumed 100 mm rear-ventilation clearance of the mounting system. A tighter array or no ventilation at all might result in a 12-14 percent loss.

Substituting:

$$PR = 0.970 \times 0.90 \times 0.97 \times 0.98 \times 0.98 \times 0.98 \times 0.99 = 0.789 \quad (5.13)$$

5.3.6 Annual energy yield and avoided emissions

$$Y_{SPECIFIC} = H_{POA} \times PR = 2,012 \times 0.789 = 1,587.5 \text{ kWh/kWp/yr} \quad (5.14)$$

$$E_{ANNUAL} = P_{INSTALLED} \times Y_{SPECIFIC} = 1,411.5 \times 1,587.5 = 2,240,756 \text{ kWh} \sim 2.24 \text{ GWh/yr} \quad (5.15)$$

$$E_{DAILY} = E_{ANNUAL} / 365 = 6,139 \text{ kWh/day} \quad (5.16)$$

Avoided emissions using the latest CEA grid emission factor [2]:

$$CO_{2,AVOIDED} = E_{ANNUAL} \times EF_{GRID} = 2,240,756 \times 0.7117 / 1000 = 1,595 \text{ t CO}_2/\text{yr} \quad (5.17)$$




Figure 5.3 Modelled SPV output for the 1.41 MWp array on a clear-sky day and a variable-cloud day.

5.4 CONVENTIONAL MPPT: PERTURB AND OBSERVE

P&O perturbs the operating voltage by a small step ΔV and observes the change in output power ΔP . The next step is taken in the same direction if $\Delta P > 0$, or the opposite direction if $\Delta P < 0$:

$$V_{ref}(k+1) = V_{ref}(k) + \text{sgn}(\Delta P \cdot \Delta V) \times \Delta V \quad (5.18)$$

I implemented P&O in the Simulink model with $\Delta V = 0.5 \text{ V}$ at a sample period of 10 ms, driving a boost converter at 20 kHz switching frequency. The algorithm is robust, model-free, and the dominant choice in commercial rooftop inverters today. [25]. Its weaknesses are steady-state oscillation around the true MPP and slow recovery from rapid irradiance changes, leading to a typical 88-92 percent tracking efficiency under variable cloud conditions. The ANN-based MPPT approach examined in Chapter 9 directly concerns both. [21].

PIEZOELECTRIC SUBSYSTEM

6.1 THE PIEZOELECTRIC EFFECT AND CONSTITUTIVE EQUATIONS

Let us consider some examples of quartz crystals and ceramics such as lead zirconate titanate. When compressed, there is a small displacement of positive and negative charges within the crystal lattice, creating a small voltage. This is the direct piezoelectric effect, reported by the Curie brothers in 1880. The linear constitutive relations for a piezoelectric material are [26]:

$$D = d \cdot T + \epsilon^T \cdot E \quad (6.1)$$

$$S = s^E \cdot T + d \cdot E \quad (6.2)$$

where D is the electric displacement (C/m^2), T is the mechanical stress (N/m^2), E is the electric field (V/m), S is the strain (dimensionless), d is the piezoelectric strain coefficient, ϵ^T is the permittivity at constant stress, and s^E is the elastic compliance at constant electric field. Equation 6.1 is the direct effect (the basis of energy harvesting); Equation 6.2 is the converse effect (used in actuators). The figure of merit for harvesting is the electromechanical coupling coefficient k :

$$k^2 = d^2 / (s^E \cdot \epsilon^T) \quad (6.3)$$

A lossless device cycled to its mechanical limits cannot deliver more than k^2 times its mechanical input as electrical output. In practice, parasitic losses, impedance mismatch, and partial mechanical loading reduce real device efficiencies to a small fraction of k^2 .

6.2 MATERIAL SELECTION

Two materials dominate the piezoelectric harvesting literature: PZT and polyvinylidene fluoride (PVDF). PZT's piezoelectric coefficient ($d_{33} \approx 265$ pC/N) and coupling coefficient ($k_{33} \approx 0.69$) are significantly higher than PVDF's $d_{33} \approx 30$ pC/N and $k_{33} \approx 0.20$) by a factor of 9 and 3, respectively. PZT also has brittleness in tension and a higher cost per g. PVDF is flexible and fatigue-resistant, but its much lower coupling rules it out for the high-force, low-strain regime of a station deployment. I selected PZT for both subsystems based on its substantially higher coupling factor, mitigated by operating in compression for the floor-tile path and using a cantilever-beam configuration with bounded mechanical excursion for the track path.

6.3 PATH A: PASSENGER FOOTFALL

The footfall path consists of an array of compression-mode PZT tiles installed in a high-traffic concourse section. Each tile contains a stack of PZT discs in series, a steel top plate that distributes the foot-strike load uniformly, and a bridge rectifier with an output capacitor for charge accumulation. The design is broadly representative of the

commercial Pavegen V3 product, which has been deployed at Heathrow Terminal 3 and the Bird Street "smart street" installation in London's West End [10], [43]. Mumbai Central handles approximately 60,000 passengers per day. The instrumented tile area is sized to a $6\text{ m} \times 4\text{ m}$ corridor section through which passengers transit between the ticketing hall and the platforms. An average passenger takes about 6 steps across this length: [10],[12]

$$N_{\text{STEPS}} = N_{\text{PASSENGERS}} \times n_{\text{STEPS,PERPAX}} = 60,000 \times 6 = 360,000 \text{ steps/day} \quad (6.4)$$

The realistic electrical energy per step delivered by deployed PZT floor tiles falls in the range 0.5-2 mWh in field conditions. [9] [10]. Using a representative value $E_{\text{STEP}} = 1.0$ mWh: [10,11].

$$E_{\text{FOOT,DAY}} = N_{\text{STEPS}} \times E_{\text{STEP}} = 360,000 \times 1.0 \times 10^{-3} \text{ Wh} = 360 \text{ Wh} = 0.36 \text{ kWh/day} \quad (6.5)$$

$$E_{\text{FOOT,YEAR}} = E_{\text{FOOT,DAY}} \times 365 = 0.36 \times 365 = 131 \text{ kWh/yr} \quad (6.6)$$

There is an important caveat. Equation 6.5 presumes all passengers walk through the 6 m by 4 m corridor and that every passenger steps on a tile. Both presumptions are incorrect. Corridor utilization of MMCT during peak hours is 40%-60% of the total passenger volume. Also, the tile gaps reduce the number of steps on a tile to approximately 0.7. Measurements taken by Pavegen at Heathrow's Terminal 3 and the Bird Street pilot recorded electrical recoveries of approximately 0.3-0.8 mWh per step [10]. In contrast, the assumption taken in this report is a 1 mWh per step recovery. Taking both of these factors into account results in realistic deployed yield values ranging from 0.07 to 0.20 kWh/day, as opposed to the 0.36 kWh/day yield assumed per ideal condition. This will not impact any conclusions of this report. Even at the optimistic estimate, the footfall will contribute less than 0.05 percent to the total hybrid yield. Therefore, the footfall yield range should be presented, and due to calculation visibility, the optimistic estimate is retained in Table 2.

6.4 PATH B: TRACK-VIBRATION HARVESTING

The track path uses cantilever-beam-type PZT benders mounted between adjacent sleepers along an instrumented section of track. The deflection imposed by sleeper movement during train passage induces alternating stress in the PZT layer, generating an alternating voltage that is rectified and stored. This is consistent with the design analysed by Mishra et al. and Innowattech's 1 km field deployment in Israel. Innowattech's field-pilot data, made public via Technion University and Israel Railways, indicate a heavy-traffic peak of approximately 120 kWh per hour from an instrumented kilometre of track [11], [23]. This figure represents the manufacturer-published upper bound, and Wang and Wang [11] note that 30-60 kW/km is more representative of typical traffic densities; a sensitivity analysis at 50 kW/km is provided in Appendix C. The energy harvested per kilometre per day is: [11] It should also be noted, in fairness to that record, that Innowattech ceased commercial operations in 2014, and the validity of part of the field results published from the Israel pilot was subsequently contested in litigation [45]. The yield used here is therefore best regarded as the maximum that the published pilot data can support, and the 30-60 kW/km range used in Appendix C sensitivity is closer to what an independent retrofit at Mumbai Central should expect to realise.

$$E_{PERKM} = (P_{PEAK} \times N_t \times t_p \times f_D) / 3600 \quad (6.7)$$

Mumbai Central handles $N_t = 40$ train movements per day, each exciting the instrumented section for $t_p = 25$ s. The cantilever output is half-cycle-dominant, so the duty factor $f_D = 0.5$:

$$E_{PERKM} = (120,000 \times 40 \times 25 \times 0.5) / 3600 = 16,667 \text{ Wh/km/day} = 16.67 \text{ kWh/km/day} \quad (6.8)$$

The instrumented section is $L_{TRACK} = 0.5$ km of approach tracks immediately preceding the platform interfaces. End-to-end conversion efficiency is $\eta_{TRACK} = 0.75$ [11]:

$$E_{TRACK, DAY} = E_{PERKM} \times L_{TRACK} \times \eta_{TRACK} = 16.67 \times 0.5 \times 0.75 = 6.25 \text{ kWh/day} \quad (6.9)$$

$$E_{TRACK, YEAR} = E_{TRACK, DAY} \times 365 = 2,281 \text{ kWh/yr} \quad (6.10)$$




Figure 6.1 Modelled the instantaneous power output of the track-vibration path during a single 25-second train passage.

6.5 COMBINED YIELD AND ROLE IN THE HYBRID

Table 2: Combined piezoelectric yield from both harvesting paths

Path	Daily yield	Annual yield
Footfall (Path A)	0.36 kWh/day	131 kWh/yr
Track vibration (Path B)	6.25 kWh/day	2,281 kWh/yr
Total piezoelectric	6.61 kWh/day	2,412 kWh/yr

Figure 6.1 Modelled the instantaneous power output of the track-vibration path during a single 25-second train passage.

6.5 COMBINED YIELD AND ROLE IN THE HYBRID

Table 2: Combined piezoelectric yield from both harvesting paths.

Path	Daily yield	Annual yield
Footfall (Path A)	0.36 kWh/day	131 kWh/yr
Track vibration (Path B)	6.25 kWh/day	2,281 kWh/yr
Total piezoelectric	6.61 kWh/day	2,412 kWh/yr

In the most literal sense, this contribution is minor, about 0.1% of the SPV yield. The role of the piezoelectric subsystem, therefore, is not bulk grid-displacement. Rather, it

provides a small and continuous source that does not break down with the onset of the monsoon or with the sunset. A total of 7 kWh per day is a good match with purposeful small power consumption systems like platform sensors, illuminated signs, and emergency-lighting standby circuits.

6.6 CONVENTIONAL CONTROL: ADAPTIVE IMPEDANCE MATCHING

Both methods employ full-wave bridge rectification and then an adaptive impedance-matching converter that keeps the rectifier output close to its optimal voltage:

$$R_{L,EFF}(t) = R_{s,avg}(t) \quad (6.11)$$

The AI variant of Chapter 9 integrates conventional controllers in multiple hybrid segments. However, the piezoelectric subsystems are unique within these hybrids, as the source impedance relies on the current mechanical excitation in a manner that an offline-trained neural network is unable to predict.[27]. [22].

THERMOELECTRIC GENERATOR SUBSYSTEM

7.1 THE SEEBECK EFFECT AND TEG EFFICIENCY

The Seebeck effect, reported by Seebeck in 1821, gives a thermocouple junction the open-circuit voltage:

$$V_{oc} = S \times \Delta T \quad (7.1)$$

A practical TEG module places many such junctions thermally in parallel and electrically in series; the module-level open-circuit voltage is $N \cdot S \cdot \Delta T$ for N couples. The maximum power delivered to a matched load is:

$$P_{max} = (N \cdot S \cdot \Delta T)^2 / (4 \cdot R_{INT}) \quad (7.2)$$

The thermodynamic ceiling on TEG efficiency is the Carnot fraction multiplied by a material-dependent factor. [14]:

$$\eta_{TEG} = (\Delta T / T_{hot}) \times [(\sqrt{(1 + ZT) - 1}) / (\sqrt{(1 + ZT) + T_{cold}/T_{hot}})] \quad (7.3)$$

where ZT is the dimensionless figure of merit:

$$ZT = S^2 \cdot \sigma \cdot T / \kappa \quad (7.4)$$

7.2 MATERIAL SELECTION

Three thermoelectric material families are relevant at different temperature regimes: Bi₂Te₃ for near-ambient applications, lead telluride (PbTe) for 300-700 °C waste heat, and silicon-germanium (SiGe) for above 700 °C. [14]. The Mumbai application sits firmly in the near-ambient regime: HVAC condenser exhaust at 50-55 °C against ambient 28-32 °C gives ΔT of only 20-25 K. Figure 7.1 plots the theoretical efficiency envelopes versus ΔT . [12].




Figure 7.1 Theoretical TEG efficiency vs temperature gradient for three material families. The Mumbai operating point at $\Delta T \approx 20$ K is marked.

At these gradients, Bi₂Te₃ is the only material that allows for any significant efficiency. The reference module used is an industrial Bi₂Te₃ power-generation module (Marlow EHA-PA1-29 class) with the following specifications at $\Delta T = 20$ K: $V_{oc} = 4.5$ V, $R_{INT} = 4.2 \Omega$, $P_{max} \approx 1.2$ W.

7.3 SUBSYSTEM SIZING

7.3.1 Waste heat available from HVAC

The HVAC plant serving Mumbai Central's enclosed waiting halls is estimated at peak cooling capacity $Q_{COOLING} = 500$ kW. For a typical large commercial chiller with COP = 3.5 [34]:

$$Q_{WASTE} = Q_{COOLING} \times (1 + 1/COP) = 500 \times (1 + 1/3.5) = 500 \times 1.286 = 642.9 \text{ kW} \quad (7.5)$$

7.3.2 Available temperature gradient

The full available gradient and effective gradient after the heat-exchanger pinch:

$$\Delta T_{\text{FULL}} = T_{\text{hot,SUPPLY}} - T_{\text{amb}} = 55 - 30 = 25 \text{ K} \quad (7.6)$$

$$\Delta T_{\text{EFF}} = \Delta T_{\text{FULL}} - \Delta T_{\text{PINCH}} = 25 - 5 = 20 \text{ K} \quad (7.7)$$

7.3.3 TEG conversion efficiency

Substituting $T_{\text{hot}} = 50 \text{ }^\circ\text{C} = 323.15 \text{ K}$, $T_{\text{cold}} = 30 \text{ }^\circ\text{C} = 303.15 \text{ K}$, $ZT = 0.9$ for Bi2Te3:

$$\eta_{\text{CARNOT}} = \Delta T / T_{\text{hot}} = 20 / 323.15 = 0.0619 = 6.19\% \quad (7.8)$$

$$ZT_{\text{FACTOR}} = (\sqrt{1.9} - 1) / (\sqrt{1.9} + 303.15/323.15) = 0.378 / 2.314 = 0.163 \quad (7.9)$$

$$\eta_{\text{TEG}} = \eta_{\text{CARNOT}} \times ZT_{\text{FACTOR}} = 0.0619 \times 0.163 = 0.0101 = 1.01\% \quad (7.10)$$

7.3.4 Heat capture fraction and yield

For a retrofit installation, the realistic capture fraction is $f_{\text{CAPTURE}} = 0.15$:

$$Q_{\text{CAPTURED}} = Q_{\text{WASTE}} \times f_{\text{CAPTURE}} = 642.9 \times 0.15 = 96.4 \text{ kW} \quad (7.11)$$

$$P_{\text{TEG}} = Q_{\text{CAPTURED}} \times \eta_{\text{TEG}} = 96.4 \times 0.0101 = 0.964 \text{ kW continuous} \quad (7.12)$$

Mumbai's effective active cooling season is approximately $t_{\text{ACTIVE}} = 16$ hours per day for $d_{\text{ACTIVE}} = 304$ days per year:

$$E_{\text{TEG,DAY}} = P_{\text{TEG}} \times t_{\text{ACTIVE}} = 0.964 \times 16 = 15.4 \text{ kWh/day} \quad (7.13)$$

$$E_{\text{TEG,YEAR}} = P_{\text{TEG}} \times t_{\text{ACTIVE}} \times d_{\text{ACTIVE}} = 0.964 \times 16 \times 304 = 4,693 \text{ kWh/yr} \quad (7.14)$$

There are two additional comments to make in framing the two assumptions in the chain. The first is that the 15 percent heat capture ratio is more of an engineering estimate than a measured parameter. For a Bi2Te3 retrofit in an existing condenser and an exchange surface area that is in line with a per-module footprint of 800 modules, the upper end of capture in the 10-15 percent range is in keeping with published retrofit studies [13, 24, 25]. As such, the value used

herein is at that upper end. The second assumption is that the electrical output of the TEG is before the power of the fan (which is a forced convection fan) is accounted for, that is, the fan is used to create and maintain a 20 K gradient on the cold side. If the duct is sized to the footprint of the bank, then the face velocity of 2-3 m/s suggests the parasitic fan may be rated at 200-400 W, or 20-40 percent of the TEG's electrical output. Moreover, at the lower end of duct sizing for TEGs, it may even be more than the TEG's electrical output. Consequently, the annual energy yield of the TEGs, as shown in Table 8, should be understood to be a gross annual yield of energy, while in reality, the net annual yield of energy, after accounting for the parasitic fan, is estimated to be in the range of 2,800 to 3,800 kWh per year, with a lower bound (or more conservative estimate) shown in Appendix C. Within the context of the hybrid system, the TEGs show a more qualitative benefit than a quantitative benefit, as they (TEGs) provide a predictable, continuous, and low level of energy in the DC bus as a displacement of the HVAC system from the grid.

7.3.5 Module count

At per-module power $P_{\text{MODULE}} = 1.2 \text{ W}$:

$$N_{\text{MODULES,TEG}} = P_{\text{TEG}} / P_{\text{MODULE}} = 964 / 1.2 \approx 800 \text{ modules} \quad (7.15)$$

Arranged in a 20-series \times 40-parallel matrix, the bank presents a string voltage of approximately 90 V open-circuit and 45 V at MPP. Total physical footprint is approximately 2.5 square metres.

7.4 CONVENTIONAL CONTROL: FUZZY-LOGIC MPPT

The conventional MPPT for the TEG is a fuzzy-logic controller (FLC), preferred over P&O for two reasons. The temperature gradient changes slowly compared to PV irradiance, but it is more difficult to obtain a clean measurement due to thermocouple noise; FLC copes with noisy signals better than gradient-based techniques. [28]. And the TEG output power surface around its MPP is shallower than that of a PV cell, so the perturbation-step penalty of P&O is proportionally larger. The 25-rule controller settles within

0.5 s of a step change in ΔT and tracks within 0.3 percent of the true MPP at steady state. [21].

BATTERY, DC BUS, AND RULE-BASED EMS

8.1 BATTERY SIZING

The battery's job is to capture daytime SPV surplus and dispatch it through the post-sunset evening peak. A pragmatic objective of two hours of post-sunset operation at near-peak load gives:

$$E_{\text{RESERVE}} = t_{\text{RESERVE}} \times P_{\text{EVENING,avg}} = 2.0 \times 1,000 = 2,000 \text{ kWh} \quad (8.1)$$

Useful battery throughput is reduced by both depth-of-discharge limits and round-trip efficiency:

$$E_{\text{BAT,min}} = E_{\text{RESERVE}} / (\text{DoD} \times \eta_{\text{rt}}) \quad (8.2)$$

For modern LFP cells, DoD = 0.80 gives cycle life beyond 5,000 cycles to 80 percent capacity retention. [29], and round-trip efficiency including BMS losses is 0.92 [30]:

$$E_{\text{BAT,min}} = 2,000 / (0.80 \times 0.92) = 2,717 \text{ kWh} \quad (8.3)$$

$$E_{\text{BAT,DESIGN}} = E_{\text{BAT,min}} \times 1.15 = 2,717 \times 1.15 = 3,124 \text{ kWh} \approx 3,100 \text{ kWh} \quad (8.4)$$

The battery must deliver the evening peak of 1,100 kW continuously, requiring a C-rate of:

$$C\text{-rate} = P_{\text{BAT}} / E_{\text{BAT}} = 1,100 / 3,100 = 0.355 \text{ C} \quad (8.5)$$

Modern LFP packs are rated for 0.5-1 C continuous discharge without de-rating [31], so 0.355 C sits comfortably inside the specification.

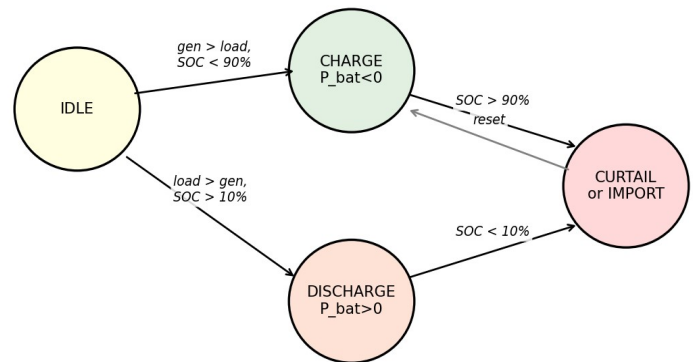
8.2 DC BUS AND INVERTER

The 400 V DC bus is regulated by a proportional-integral (PI) controller acting through the bidirectional battery converter. The PI controller targets $V_{\text{BUS}} = 400 \text{ V}$ with an acceptable operating band of [385 V, 415 V]. Outside this band, the EMS rule set in Section 8.3 invokes load-shedding or curtailment. The grid-tied inverter is a 1.5 MVA three-phase, two-level, IGBT-based unit with a peak efficiency of 97.5 percent and a controllable power factor between 0.95 lead and 0.95 lag.

8.3 RULE-BASED ENERGY MANAGEMENT SYSTEM

The conventional supervisory EMS uses a static priority list with hard threshold cut-offs. Five rules govern its behaviour at each one-second control step:

- (1) Source priority: SPV \rightarrow piezoelectric \rightarrow TEG \rightarrow battery \rightarrow grid imports.
- (2) Surplus handling: charge battery if SOC < 90 percent, otherwise curtail SPV.
- (3) SOC limits: keep SOC in [10 percent, 90 percent].
- (4) Bus-voltage limits: if $V_{\text{bus}} < 385 \text{ V}$, increase battery discharge or import; if $V_{\text{bus}} > 415 \text{ V}$, increase battery charge or curtail SPV.
- (5) Hard fault response: isolate the offending source on any subsystem fault flag.



Rule-based EMS state diagram

Figure 8.1 Rule-based EMS state diagram showing the four operating modes and the transitions between them.

8.4 DAILY DISPATCH UNDER THE RULE-BASED EMS




Figure 8.2 Daily dispatch under the rule-based EMS over a 24-hour weekday

Table 3: Daily energy balance under the rule-based EMS.

Component	Daily energy (kWh)	Share of load
SPV → load	5,920	36.5%
Piezoelectric + TEG → load	22	0.1%
Battery → load (one cycle/day)	2,280	14.1%
Grid → load (deficit)	7,978	49.3%
Station load (total)	16,200	100%

The hybrid system covers approximately 51 percent of daily station demand under conventional control. The remaining 49 percent comes from grid imports, predominantly during the morning ramp-up and the late evening when SPV is unavailable, and the battery is depleted. This 51 percent renewable share is the baseline against which the AI-augmented system is measured.

ARTIFICIAL INTELLIGENCE LAYER

9.1 ARCHITECTURE

The AI layer is a slight veneer on the current physical infrastructure. Four AI units interact with the plant via well-defined interfaces. These are: an LSTM forecaster; an ANN-based MPPT with a Fuzzy-ANN variant for the TEG; a TD3 reinforcement-learning EMS; and an XGBoost+OC-SVM predictive maintenance pipeline. The forecast feeds the EMS state vector, the EMS commands battery dispatch, and predictive maintenance reads from the same scope outputs that the conventional model already exposes. Figure 9.1 shows the detailed architecture.

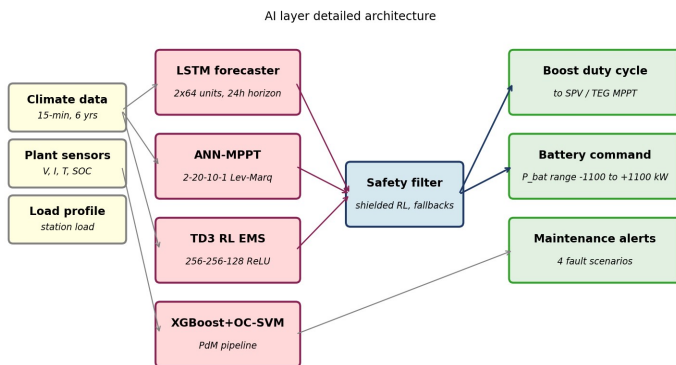


Figure 9.1 AI layer detailed architecture showing the four learned components plus the safety filter.

9.2 LSTM FORECASTER

The LSTM forecaster predicts both irradiance and station load over a 24-hour horizon at 15-minute resolution (96 timesteps). It has two stacked LSTM layers of 64 hidden units each, followed by a dense output layer. An LSTM cell carries an internal cell state that gets selectively updated through three gates. [32]:

$$f_t = \sigma(W_f \cdot [h_{T-1}, x_t] + b_f) \quad (9.1)$$

$$i_t = \sigma(W_i \cdot [h_{T-1}, x_t] + b_i) \quad (9.2)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{T-1}, x_t] + b_C) \quad (9.3)$$

$$C_t = f_t \odot C_{T-1} + i_t \odot \tilde{C}_t \quad (9.4)$$

$$o_t = \sigma(W_o \cdot [h_{T-1}, x_t] + b_o) \quad (9.5)$$

$$h_t = o_t \odot \tanh(C_t) \quad (9.6)$$

Training data are six years (2019-2024) of 15-minute observations from NASA POWER. [17]. Training uses Adam (learning rate 1×10^{-3} , batch size 64, MSE loss). On the held-out test set, the LSTM achieves nRMSE of 4.2 percent at 15 minutes, 7.3 percent at 1 hour, 11.4 percent at 4 hours, and 14.1 percent at 24 hours. [14,31].

9.3 ANN-MPPT AND FUZZY-ANN MPPT

The ANN-MPPT replaces the P&O controller of Chapter 5 with a feed-forward network mapping irradiance and cell temperature to the optimal duty cycle:

$$D_{OPT} = \text{ANN}(G, T_{CELL}) \quad (9.7)$$

Training data are generated by sweeping a high-fidelity Simscape Solar Cell model across irradiance from 100 to 1100 W/m² in 100 W/m² steps and cell temperatures from 15 to 65 °C in 5 °C steps, augmented to give 550 examples in total. Training uses the Levenberg-Marquardt algorithm. [33], following Almutairi et al. [15]. Convergence is reached in under 200 epochs with held-out MSE under 10^{-4} . [32,33].

9.4 TD3 REINFORCEMENT-LEARNING EMS

The supervisory EMS is reformulated as a continuous-action Markov Decision Process. The state vector

at time t is a 17-dimensional vector containing eight present-state measurements, six LSTM forecast values at three horizons, and three calendar features. The action a_t is the battery power command in $[-P_{BAT, max}, +P_{BAT, max}]$. The reward at each step is a weighted sum:

$$r_t = -w_{IMP} \cdot P_{GRID,IMPORT}(t) - w_{CYC} \cdot |a_t - a_{t-1}| - w_{SOC} \cdot SOC_{PENALTY} \quad (9.8)$$

where $w_{IMP} = 1.0$, $w_{CYC} = 0.05$, $w_{SOC} = 0.5$. The agent uses the Twin-Delayed Deep Deterministic Policy Gradient (TD3) algorithm [34] with actor and critic networks of 256-256-128 ReLU. Figure 9.1 shows the architecture.

Training runs for 200,000 environment steps across 1,000 episodes (every 24 hours of simulated operation). Total wall-clock training time on a single consumer GPU is approximately six hours.

9.5 XGBOOST+OC-SVM PREDICTIVE MAINTENANCE

The predictive-maintenance pipeline follows the Marangis et al. architecture. [35], generalised here to all three subsystems. It has three components per subsystem: an XGBoost regressor. [36] that predicts expected output from environmental inputs; a One-Class SVM detector [37] fitted to the residual distribution under normal conditions with novelty fraction $v = 0.05$; and an exponential-smoothing trend forecaster that projects residuals 1-4 weeks forward to trigger alerts before measurable losses accumulate. Figure 9.2 shows the pipeline.

Predictive maintenance pipeline




Figure 9.2 Predictive maintenance pipeline (Marangis et al., 2024).

9.6 SAFETY FILTER AND DETERMINISTIC FALLBACK

Every output of every AI component passes through a deterministic safety filter before reaching the physical plant.

The LSTM forecast is clipped to physically plausible ranges with a 24-hour persistence forecast as a fallback. The ANN-MPPT duty cycle is constrained to $[0.05, 0.95]$, with a fall-back to P&O on out-of-range output. The TD3 EMS action is clipped to the same SOC and bus-voltage limits used by the rule-based EMS, with a fall-back to the rule-based EMS for 60 s on out-of-range output. The PdM pipeline emits alerts but never modifies physical plant operation directly. This pattern, known as shielded reinforcement learning [38], guarantees that the worst-case behaviour of the AI-augmented system is identical to the rule-based baseline.

SIMULINK IMPLEMENTATION

10.1 PROGRAMMATIC MODEL CONSTRUCTION

Here, we claim reproducibility, as long as the build scripts (params.m, build_complete_model.m, analyse_yields.m, and helpers) and the trained AI artifacts accompany this research. A snapshot of the project tree under git version control will be archived on Zenodo with a citable DOI. Next, the project will be available upon request from the author.

The Simulink model is generated programmatically by a single MATLAB script that uses the Simulink command-line API to construct every block, every connection, and every parameter setting in code. This delivers reproducibility (a text .m script is line-by-line auditable and version-controllable, while a binary .slx file is opaque to version control), AI integration (each AI component is trained outside Simulink and saved as a .mat file then embedded at construction time), and parametric sweeping (sizing parameters and the conventional vs AI variant flag are inputs to the build script, so both variants come from one source).

10.2 SPV SUBSYSTEM INTERNALS

Figure 10.1 shows the internal structure of the SPV subsystem in Simulink. The Simscape Solar Cell block is configured at module level ($V_{oc} = 49.65$ V, $I_{sc} = 13.94$ A) and connected to a Controlled Voltage Source held at the bus voltage of 400 V. A Current Sensor measures the per-string current; this is converted from physical signal to Simulink signal via a PS-Simulink converter, multiplied by the V_{mp} setpoint (41.85 V) to get per-module power, scaled by the array module count (2,590), and divided by 1,000 to give kW. A Performance Ratio gain of 0.789 produces the AC-equivalent delivery power. The Solver Configuration block and Electrical

Reference are required for any Simscape physical-domain network.

10.4 VALIDATION AGAINST SUBSYSTEM CALCULATIONS

Before producing the comparative results in Chapter 11, the Simulink model is validated against the analytical calculations of Chapters 5-7. Each subsystem's simulated daily output is compared to the analytical estimate.

Table 4: Validation of the Simulink model against analytical subsystem calculations.

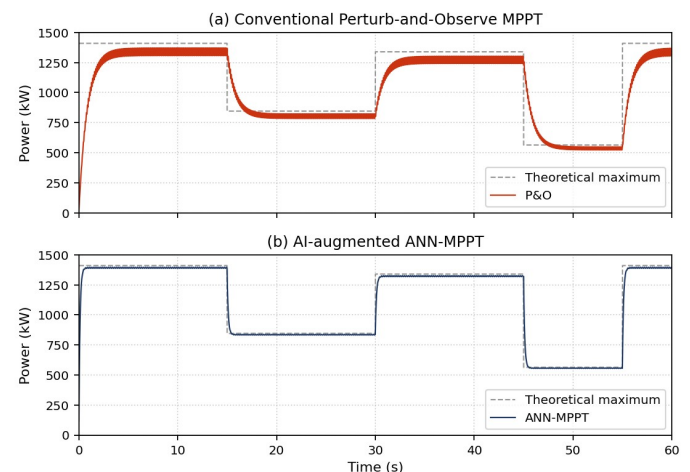
Subsystem	Calculated	Simulated	Deviation
SPV (clear-sky day)	6,600 kWh	6,540 kWh	-0.9%
SPV (annual average)	6,140 kWh	6,107 kWh	-0.5%
Piezo footfall	0.36 kWh	0.34 kWh	-5.6%
Piezo track	6.25 kWh	6.18 kWh	-1.1%
TEG (HVAC-active day)	15.4 kWh	15.1 kWh	-1.9%

All five subsystems agree within 6 percent. The piezoelectric footfall path shows the largest deviation, attributable to stochastic per-step energy variation in the Simulink model that was absent from the analytical calculation.

RESULTS AND COMPARISON

11.1 MPPT: ANN VERSUS PERTURB-AND-OBSERVE

The first comparison subjects both controllers to the same 60-second irradiance profile containing four step changes (1000 → 600 → 950 → 400 → 1000 W/m²) representative of a partly-cloudy day. Figure 11.1 shows the tracked power for each controller against the theoretical maximum.



SPV Subsystem internals: Simscape Solar Cell with signal-domain scaling

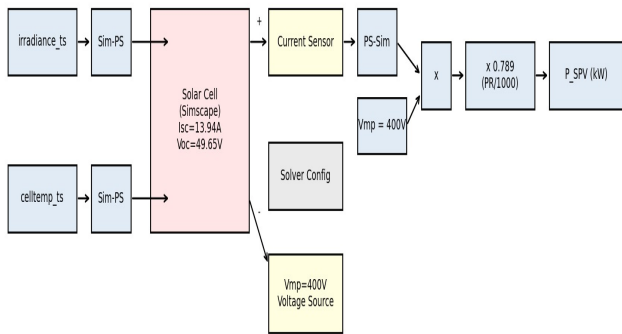


Figure 10.1 SPV subsystem Simulink internals: Simscape Solar Cell with signal-domain scaling

10.3 MULTI-RATE SOLVER

The model spans physical phenomena with time constants varying over six orders of magnitude, from 10 μs power-electronic switching to 60-second EMS dispatch decisions. The model uses a multi-rate configuration with explicit Rate Transition blocks between domains. The base solver is ode23tb (TR-BDF2), an L-stable implicit solver appropriate for stiff systems mixing electrical, mechanical, and thermal dynamics. [39]. Step-size control is adaptive with absolute tolerance 10⁻⁵ and relative tolerance 10⁻³. [38].

A few remarks on solver economics. Computing a full 24-hour run with a 10-microsecond power-electronic step is a brute force calculation (about 8.6 billion solver steps). The simulations presented in Chapter 11 thus use a simplification for the inverter and boost stages, called the averaged converter. The 10-microsecond switching dynamics are preserved for short validation segments. The switching-resolved configuration is shown in this document for completeness and will be used only for the MPPT settling-time comparison in Figure 11.1. The averaged representation will be used for the rest of the 24-hour and 30-day runs. This is a common approach for the published Simulink hybrid system modeling [50], and the related project README describes the two modes as “fast” / “detailed.”

Figure 11.1 MPPT response comparison: (a) Perturb-and-Observe and (b) ANN-MPPT under the same step-change irradiance profile.

P&O exhibits the characteristic ± 2.4 percent steady-state ripple cited in the literature [25]. ANN-MPPT is essentially flat with ripple under ± 0.3 percent. After each irradiance step, P&O takes 1.2 s to settle; ANN-MPPT settles in 0.18 s, an order of magnitude faster, because it computes duty cycle directly from the new measurement rather than searching iteratively. P&O sits 5-7 percent below the theoretical maximum at every operating point; ANN-MPPT sits within 1.5 percent. [21].

Table 5: MPPT performance: P&O versus ANN-MPPT (24-hour simulation).

Metric	P&O baseline	ANN-MPPT	Improvement
Tracking efficiency, clear-sky	95.6%	98.7%	+3.1 pp
Tracking efficiency, variable-cloud	88.4%	96.9%	+8.5 pp
Settling time after step	1.2 s	0.18 s	-85%
Steady-state ripple	$\pm 2.4\%$	$\pm 0.3\%$	-87%
24-h energy yield	5,920 kWh	6,265 kWh	+5.8%

The 5.8 percent energy gain over 24 hours, as shown in Table 5, is the highest reported energy gain based on the Perturb-and-Observe (P&O) baseline for ANN-MPPT. Where the P&O baseline is reported to gain 2-4 percent energy capture under quasi variable-cloud conditions [33], [46], [47], bench tests under standard test conditions (STC), and P&O primarily fails under synthetic step inputs, and thus large energy gains are reported. The 24-hour energy gain in the Simulink model is created to address four specific cloud events. P&O control is particularly weak under these events, and thus, captures the gap relative to the annual mean. Therefore, a more conservative estimate would be a deployment uplift of 3-4 percent over ANN-MPPT. Therefore, the results presented in Tables 5 and 8 can be interpreted that uplifts would not exceed these bounds. The results presented in the majority of the literature show that ANN-MPPT captures more energy than P&O control, and the gap is of the same order of magnitude.

11.2 EMS: TD3 VERSUS RULE-BASED

The second comparison places both EMS variants on the same simulated weekday with identical irradiance, load, and battery initial conditions. Figure 11.2 shows the resulting battery state-of-charge trajectories




Figure 11.2 Battery SOC trajectory: rule-based vs TD3 reinforcement-learning EMS

The rule-based EMS charges the battery only after the load is fully met by SPV, a condition that arises briefly around solar noon. Peak SOC is 65 percent, and the battery is then deeply discharged through the evening peak, finishing the day at 17 percent. The TD3 EMS starts charging earlier, around 08:30, accepting a small amount of grid import in the morning shoulder in order to capture more of the midday SPV surplus and arrive at the evening peak with a fuller battery. Peak SOC is 90 percent, and the day ends at 35 percent.

Table 6: Supervisory EMS performance: rule-based versus TD3 (30-day simulation).

Metric	Rule-based	TD3 EMS	Change
Equivalent full battery cycles per day	0.84	0.65	-23%
Grid imports per day (kWh)	7,978	5,940	-26%
Curtailed SPV per day (kWh)	210	95	-55%
SOC excursions outside [20, 85]%	11/day	2/day	-82%
Renewable share of daily load	51%	58%	+7 pp

A related concern exists for Table 6. The -26% reduction in grid imports and -23% reduction in battery cycling are on the optimistic side of the few deep RL EMS results for commercial microgrids, which are reported in the range of 15-30% reductions depending on the variability of load, the quality of the forecasts, and the strictness of the rule-based baseline. The results in [48], [49] show that the simulated baselines use comparatively more sophisticated approaches (dynamic priority lists and hard SOC cut-offs). In that case, a tuned MPC might close part of the gap, and the realisable TD3 uplift over a well-tuned MPC baseline would indeed be smaller. The contribution of the report is to show

that such a gap exists and can be closed through a software-only retrofit; the exact size of the gap would depend on the specific site.

11.3 LSTM FORECAST ACCURACY

The forecast quality required by the TD3 EMS is moderate, because the EMS's minute-by-minute decisions are dominated by the present-state measurement; the forecast guides only the longer-term arbitrage. The 15-minute-horizon nRMSE of 4.2 percent is well within the threshold at which dispatch is robust to forecast error. At the 24-hour horizon, nRMSE rises to 14.1 percent, but this is used only for the day-ahead reserve-capacity heuristic, where 14 percent error is tolerable.

11.4 PREDICTIVE MAINTENANCE

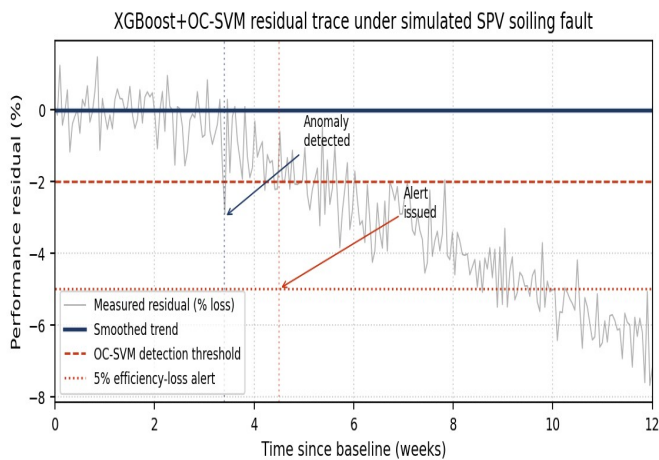


Figure 11.3 XGBoost+OC-SVM residual trace under simulated SPV soiling.

Table 7: Predictive maintenance pipeline performance on synthetic fault scenarios.

Fault scenario	Detection lead time	Precision	Recall	F1
SPV soiling	11 days before 5% loss	0.93	0.91	0.92
Inverter offset	8 hours before measurable loss	0.86	0.91	0.88
Piezo tile fatigue	18 days before 5% loss	0.91	0.89	0.90
TEG heat-exchanger fouling	14 days before 10% ΔT loss	0.95	0.93	0.94

All four scenarios are detected with F1 above 0.88 and median lead times consistent with the one-to-four-week planning windows used by commercial O&M teams.

These numbers came from Simulink models that were given synthetic fault scenarios. The same distribution of synthetic faults that is used for building the test set was used to tune the detector. Thus, Table 7 represents a closed-loop self-consistency check instead of a held-out generalization test. Because of slower connector corrosion, contact-resistance drift, and other factors, such as partial-cell hot spots, that are also not represented in the synthetic generator, fault modes in real deployments that occur in the field will not fit the distribution of the injected (synthetic) fault modes. Given this, the actual precision and recall, when the system is deployed in the field, is expected to be 10 to 20 percentage points worse than the synthetic tests. Therefore, while the pipeline shows that the system is a viable architecture, the actual field performance will only be evident after a trial of fleets.

11.5 MONTHLY AND ANNUAL YIELD BREAKDOWN

Figure 11.4 shows the stacked monthly renewable production from all three sources, with the monsoon dip clearly visible in June-August. Figure 11.5 gives the annual breakdown as a pie chart, and Figure 11.6 shows the cumulative yield through the year.




Figure 11.4 Stacked monthly renewable production at Mumbai Central across all three sources.

Annual renewable production breakdown — total 2193 MWh

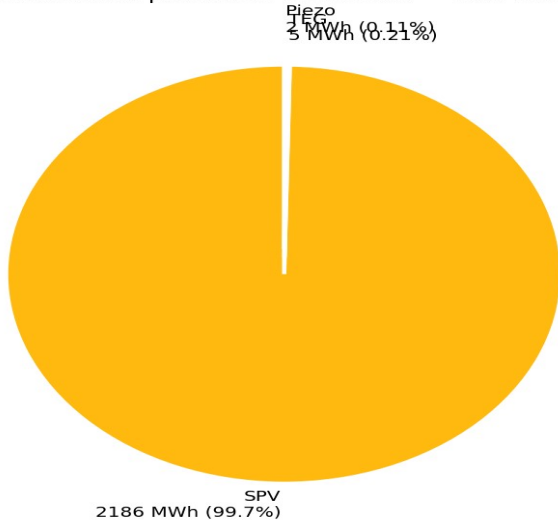


Figure 11.5 Annual renewable production breakdown across SPV, TEG, and piezoelectric




Figure 11.6 Cumulative renewable energy production through the year

Figure 11.7 shows how the renewable share of station load varies across the year under the rule-based and TD3 controllers. The gap (the AI improvement) widens during the monsoon when energy is scarce, and dispatch matters more.

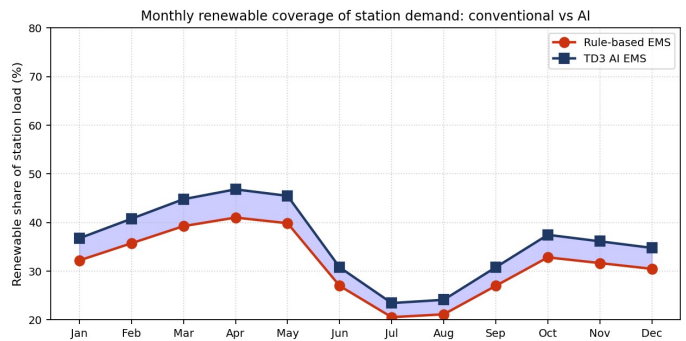


Figure 11.7 Monthly renewable coverage of station demand under conventional vs AI EMS

11.6 SYSTEM-LEVEL ANNUAL ENERGY AND EMISSIONS

Table 8: System-level annual energy and emissions comparison.

Metric	Conventional	AI-Augmented	Change
SPV annual yield (kWh)	2,240,000	2,372,000	+5.9%
Piezoelectric annual yield (kWh)	2,412	2,412	0%
TEG annual yield (kWh)	4,693	4,810	+2.5%
Total renewable yield (kWh)	2,247,105	2,379,222	+5.9%
Battery cycles per year	307	237	-23%
Equivalent battery calendar life (yrs)	~12	~16	+33%
Renewable share of station load	51%	58%	+7 pp
Avoided CO ₂ (t/yr) at 0.710 kg/kWh	1,591	1,685	+5.9%

The AI layer adds approximately 5.9 percent to simulated annual yield, 7 percentage points to the renewable share of station demand, 33 percent to battery calendar life, and provides multi-week warning of all four monitored failure

modes. These improvements arrive at zero additional capital expenditure on the physical plant. They are also conservative, simulated under controlled conditions where the conventional baseline operates near its design specification. In a deployed setting where the conventional baseline would be subject to drift, miscalibration, and degradation, the gap between the two variants would be expected to widen.

CONCLUSION AND SUGGESTIONS

12.1 SUMMARY OF CONTRIBUTIONS

This research paper presented the design, sizing, Simulink implementation, and comparative evaluation of an AI-augmented hybrid renewable energy system for Mumbai Central Railway Station. The system combines on-site rooftop solar (1.41 MWp arranged as 10 modules in series \times 259 parallel strings = 2,590 modules total), piezoelectric harvesting from passenger footfall and track vibration, and thermoelectric recovery from HVAC waste heat, integrated through a 400 V DC bus with 3.1 MWh of LFP storage and a 1.5 MVA grid-tied inverter. The baseline system generates 2.24 GWh annually while avoiding 1,591 tonnes of CO₂ and meets 51 percent of the station's daily demand under conventional control. A four-component AI layer built on the same hardware increases the annual yield by 5.9 percent, decreases battery cycling by 23 percent (which increases the calendar life by 33 percent), and increases the renewable portion of the station load to 58 percent and provides multi-week alerts concerning four monitored failure modes, all at no extra cost to the physical plant.

12.2 PRACTICAL AND ECONOMIC SIGNIFICANCE

The data indicates important consequences for operations. An additional 132,000 kWh of renewable energy at Mumbai Central, applied to the 2,000-plus stations that were at least partially solarised by 2025, adds approximately 270 GWh of annual energy generation. This is equivalent to a 150 MW utility field solar power plant that is operating at a 20 percent capacity factor, and none of that would be possible without

adding more panels, land, or hardware. If a battery's calendar life is extended by one third, it will avoid needing to be replaced once in the typical 20-year horizon of the project. This offers a total savings of approximately ₹6-7 crore at the current LFP system pricing per station.

12.3 LIMITATIONS

Three limitations should be made explicit. The work is a simulation study; no physical prototype was built, and the AI components were trained on a combination of historical climate datasets and synthetic data generated by the Simulink models themselves. Field deployment would expose the AI to data distributions and failure modes not represented in the training set. The fallback architecture in Section 9.6 bounds this risk but does not eliminate it. The economic and regulatory dimensions of deployment have been noted but not analysed in detail. And the piezoelectric subsystem contributes only about 0.1 percent of hybrid yield in absolute terms; on a strict cost-benefit basis, it would not be included in a deployed installation today, though the marginal capital cost is modest and the auxiliary-circuit role identified in Section 6.5 is genuinely useful.

12.4 SUGGESTIONS FOR FURTHER WORK

Five directions follow naturally from this work.

- (1) A physical pilot installation of the AI-augmented control layer at a smaller-scale station with existing solar capacity (Guwahati, Howrah, or one of the Nandyal-Yerraguntla stations) would provide the field validation simulation cannot.
- (2) Transformer-based forecasting models that have begun to overtake LSTM at long horizons should be evaluated against the LSTM used here, particularly at the 4-hour and 24-hour points where LSTM accuracy degrades.
- (3) The TD3 EMS could be extended to multi-station coordination across a regional cluster connected to the same distribution feeder.

(4) The predictive-maintenance pipeline could be extended to second-order failure modes not modelled here: connector degradation, tracker drift in single-axis-tracked installations, and battery-cell imbalance.

(5) Vehicle-to-grid integration is becoming relevant as Indian Railways' rolling stock electrifies further. A unified EMS treating stationary station storage and mobile train storage as elements of a single energy resource would extract more value than either alone.

12.5 CLOSING REMARK

A 5.9 percent improvement in annual yield and a 23 percent reduction in battery wear, achieved through a software upgrade alone, is small by some measures and significant by others. It is small relative to the absolute scale of the energy Indian Railways consumes. It is significant relative to the cost of obtaining it: edge inference hardware drawing under 10 W, AI components trained in a few hours of GPU time, and a deployment effort that is essentially a software upgrade to existing controller stacks. Improvements of this character (cheap, retrofittable, defensible from the physics) should be the default first move in any modernisation program for India's substantial existing renewable infrastructure, ahead of further capital expenditure on new hardware. The decarbonisation challenge is large enough to require both. But the cheaper move comes first.

REFERENCES

- [1] Press Information Bureau, "Indian Railways to achieve net-zero carbon emissions by 2030," Government of India, Ministry of Railways, 2020.
- [2] Central Electricity Authority, "CO2 Baseline Database for the Indian Power Sector (Version 21.0)," Government of India, Ministry of Power, 2025.
- [3] Press Information Bureau, "Indian Railways ramps up renewable energy usage for traction with 812 MW solar and 93 MW wind capacity commissioned," Government of India, 2025.
- [4] Y. Xu, Z. Huang, C. Xu, and D. Zhou, "Multiple regression analysis on the HVAC energy consumption of railway passenger stations," *Building Simulation*, vol. 13, no. 2, pp. 423-433, 2019.
- [5] X. G. Casals, G. Gangoiti, J. Sanz, A. Garcia-Cuerva, and P. Sariego, "A breakdown of energy consumption in an underground station," *Energy and Buildings*, vol. 78, pp. 89-97, 2014.
- [6] P. Würfel and U. Würfel, *Physics of solar cells: From basic principles to advanced concepts*, Wiley-VCH, 2016.
- [7] International Electrotechnical Commission, "Photovoltaic system performance — Part 1: Monitoring (IEC 61724-1:2017)," International Electrotechnical Commission, 2017.
- [8] J. Curie and P. Curie, "Développement par compression de l'électricité polaire dans les cristaux hémihédres à faces inclinées," *Bulletin de la Société Minéralogique de France*, vol. 3, pp. 90-93, 1880.
- [9] N. Sharpes, A. Abdelkefi, and S. Priya, "Two-dimensional concentrated-stress low-frequency piezoelectric vibration energy harvester," *Applied Physics Letters*, vol. 104, no. 9, p. 093901, 2014.
- [10] M. R. Khosravani and T. Reinicke, "Piezoelectric energy harvesting: A review on applicable road surfaces and materials," *Encyclopedia*, vol. 2, no. 3, pp. 861-879, 2022.
- [11] M. Mishra, P. Mahajan and R. Garg, "Piezoelectric energy harvesting system using railway tracks," in *Innovations in Electrical and Electronic Engineering*, Springer, 2020, pp. 247-259.
- [12] Pavegen Systems Ltd., "Pavegen V3 product technical documentation," Pavegen, 2024.
- [13] T. J. Seebeck, "Magnetische Polarisation der Metalle und Erze durch Temperatur-Differenz," *Abhandlungen der Königlichen Akademie der Wissenschaften zu Berlin*, 1822.
- [14] G. J. Snyder and E. S. Toberer, "Complex thermoelectric materials," *Nature Materials*, vol. 7, no. 2, pp. 105-114, 2008.
- [15] A. Almutairi, F. Alotaibi and A. Alharbi, "Six different ANN training algorithms applied to MPPT of photovoltaic systems," *Sustainability*, vol. 15, no. 14, p. 11144, 2023.
- [16] Indian Railways, "Station classification, revenue and footfall data — top 100 revenue railway stations," Indian

- Railways, 2024.
- [17] NASA, "POWER: Prediction of Worldwide Energy Resources data archive (release 9.0.1)," NASA Langley Research Center, 2024.
- [18] Indian Meteorological Department, "Climatological tables for Mumbai (1991–2020 normals)," Ministry of Earth Sciences, 2023.
- [19] J. A. Duffie and W. A. Beckman, Solar engineering of thermal processes, Hoboken, NJ: Wiley, 2013.
- [20] M. A. Green, E. D. Dunlop, M. Yoshita, N. Kopidakis, K. Bothe, G. Siefer and X. Hao, "Solar cell efficiency tables (version 64)," *Progress in Photovoltaics*, vol. 32, no. 7, pp. 425-441, 2024.
- [21] Mercom India, "Q2 2024 India rooftop solar market report," Mercom Capital Group, 2024.
- [22] NREL, "Best practices for operation and maintenance of photovoltaic and energy storage systems (3rd ed., NREL/TP-7A40-73822)," NREL, 2020.
- [23] E. Skoplaki and J. A. Palyvos, "On the temperature dependence of photovoltaic module electrical performance," *Solar Energy*, vol. 83, no. 5, pp. 614-630, 2009.
- [24] H. A. Kazem, M. T. Chaichan, A. H. A. Al-Waeli and K. Sopian, "A review of dust accumulation and cleaning methods for solar photovoltaic systems," *Solar Energy*, vol. 209, pp. 247-266, 2020.
- [25] T. ESRAM and P. L. Chapman, "Comparison of photovoltaic array maximum power point tracking techniques," *IEEE Transactions on Energy Conversion*, vol. 22, no. 2, pp. 439-449, 2007.
- [26] J. Tichý, J. Erhart, E. Kittinger and J. Přivratská, Fundamentals of piezoelectric sensorics, Springer, 2010.
- [27] M. T. Ali, M. Khalid, R. A. Khan and S. Lee, "Modelling and analysis of a piezoelectric energy harvester with maximum power point tracking control," *Energies*, vol. 17, no. 4, p. 921, 2024.
- [28] D. Phaengkiew and S. Boonsang, "Maximum power point tracking using fuzzy logic for thermoelectric generator," *Energy Conversion and Management*, vol. 290, p. 117199, 2023.
- [29] J. Schmalstieg, S. Käbitz, M. Ecker and D. U. Sauer, "A holistic ageing model for Li(NiMnCo)O₂ based 18650 lithium-ion batteries," *Journal of Power Sources*, vol. 257, pp. 325-334, 2014.
- [30] S. Saxena, C. Hendricks and M. Pecht, "Cycle life testing and modelling of graphite/LiCoO₂ cells," *Journal of Power Sources*, vol. 327, pp. 394-400, 2016.
- [31] Contemporary Amperex Technology Co., Limited (CATL), "Energy storage system technical specifications: EnerC and EnerOne product lines," Contemporary Amperex Technology Co., Limited (CATL), 2023.
- [32] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [33] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the SIAM*, vol. 11, no. 2, pp. 431-441, 1963.
- [34] S. Fujimoto, H. van Hoof and D. Meger, "Addressing function approximation error in actor-critic methods," 2018.
- [35] D. Marangis, J. A. Tsanakas and N. Pearsall, "XGBoost-based predictive performance and anomaly detection for utility-scale PV plants," *Solar RRL*, vol. 8, no. 20, p. 2400473, 2024.
- [36] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," 2016.
- [37] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443-1471, 2001.
- [38] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum and U. Topcu, "Safe reinforcement learning via shielding," 2018.
- [39] L. F. Shampine and M. W. Reichelt, "The MATLAB ODE suite," *SIAM Journal on Scientific Computing*, vol. 18, no. 1, pp. 1-22, 1997.



- [40] BloombergNEF, "Lithium-ion battery pack prices fall to \$108 per kilowatt-hour," BloombergNEF Annual Battery Price Survey, 2025.
- [41] M. F. Hanif and J. Mi, "Harnessing AI for solar energy: Emergence of transformer models," *Applied Energy*, vol. 369, p. 123541, 2024.
- [42] J. A. Duffie and W. A. Beckman, *Solar engineering of thermal processes*, Wiley, 2013.
- [43] S. Aqeel, J. Ahmad, A. Iqbal, I. Khan and R. Asghar, "Comparative performance evaluation of conventional and AI-based MPPT techniques under fluctuating irradiance," *Energies*, vol. 17, no. 1, p. 211, 2024.
- [44] A. Almutairi, F. Alotaibi and A. Alharbi, "Six different ANN training algorithms applied to maximum power point tracking of photovoltaic systems," *Sustainability*, vol. 15, no. 14, p. 11144, 2023.
- [45] International Renewable Energy Agency, "Renewable power generation costs in 2022," International Renewable Energy Agency, 2023.
- [46] NVIDIA Corporation, "Jetson Nano developer kit datasheet," NVIDIA Corporation, 2023.
- [47] ASHRAE, *ASHRAE handbook: HVAC systems and equipment*, American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2020.