energies

MDPI

*Article*

# A Digital Engineering Framework for Undervoltage Risk Screening in Renewable-Rich Smart-City Distribution Networks

**Rajitha Wattegama [1,*], Michael Short [1,*], Geetika Aggarwal [1], Maher Al-Greer [1] and Raj Naidoo [2]**

[1] School of Computing, Engineering and Digital Technologies, Teesside University, Middlesbrough TS1 3BX, UK; g.aggarwal@tees.ac.uk (G.A.); m.al-greer@tees.ac.uk (M.A.-G.)
[2] Department of Electrical, Electronic and Computer Engineering, University of Pretoria, P.O. Box 14679, Hatfield, Pretoria 0028, South Africa; raj.naidoo@up.ac.za
* Correspondence: r.wattegama@tees.ac.uk (R.W.); m.short@tees.ac.uk (M.S.)

**Abstract**

Deeper penetration of renewable energy sources is seen as a vital component in the transition toward smarter, cleaner, and more intelligent cities, but injection of renewable energy sources directly into the electrical distribution network can oftentimes have unintended impacts, such as over- or undervoltage. This study explores and evaluates a digital tool for steady-state voltage screening for renewable distribution planning. Benchmark scenarios informed by South African operational records and modelled photovoltaic profiles are used, along with IEEE 33- and 69-bus feeder models, for realistic validation purposes. Eskom indicators and Tshwane/Hatfield load-shedding histories condition 10,000 test scenarios for each of the busses. AC power flow labels minimum voltages below 0.95 p.u.; all 20,000 cases converged. Ridge-logistic regression using seven operational features achieved mean precision–recall areas (PR-AUCs) of 0.9782 and 0.9476 over repeated holdouts. Applying the IEEE-33 model to IEEE-69 achieved PR-AUC 0.9376 and classification recall 0.9553, demonstrating model reuse within a shared scenario-modelling framework rather than independent-distribution generalisation. At a 30% retained-case budget, operational rankings captured 96.58%, 96.84%, and 95.66% of undervoltage cases under IEEE-33 out-of-fold, IEEE-69 out-of-fold, and model-reuse evaluation, respectively. Scenario descriptors only achieved the highest capture, 99.85%, conditional on known scenario-construction inputs. These descriptor-based results are not estimates of screening performance on real feeders; neither feature representation has field validation. The retrospective rankings select 70% fewer cases for subsequent assessment without demonstrating end-to-end acceleration, dynamic stability, or upper-voltage compliance. Urban and smart-city applications remain prospective but positive.

**Keywords:** digital engineering; renewable energy integration; undervoltage screening; distribution networks; load shedding; scenario prioritisation; physics-labelled machine learning

Check for updates

Academic Editor: Yang Du

Received: 9 October 2026
Revised: 6 November 2026
Accepted: 18 November 2026
Published: 20 November 2026

**Copyright:** © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

## 1. Introduction

Renewable and clean energy sources embedded deep with distribution networks are seen as a vital component in the transition toward smarter, cleaner, and more intelligent cities. However, the injection of renewable energy sources directly into electrical distribution network nodes can oftentimes have unintended impacts, such as over- or undervoltage, and city planners must take these impacts into account. In a highly dynamic environment, distribution network planners must evaluate changing demand, variable local generation

across multiple node locations, and changing network conditions to identify where detailed analysis is most needed. Urban energy systems often supply critical services and flexible loads, and provision for mass interconnection of PV and Wind installations, making the assessment complex beyond a single maximum-demand point. Historic evidence from known European issues involving unplanned PV expansion in Germany indicates that a lack of planning and foresight can lead to wide-scale stability problems and power outages (the '50.2 Hz problem') and subsequent wide-scale re-engineering of local renewable controls [1]. It involves uncertainty in net demand, renewable variation, and supply interruptions. Effective assessment requires indicators with clear engineering meaning and visible consequences of missed cases [2]. A voltage-magnitude violation at a solved operating point and a stability assessment answer different questions. Farrokhabadi et al. have defined and classified microgrid stability and discussed its dependence on system dynamics and disturbances [3]. Ning et al. investigated dynamic stabilisation of multimodal hybrid power supplies subjected to stepped constant-power loads using synergetic control [4]. The present work contains neither a dynamic controller model nor a stability-margin calculation. Its physical target is a steady-state undervoltage event, and its statistical target is the prioritisation of such events for subsequent electrical assessment. Consequently, the reported classification results are not evidence of voltage-collapse margins or dynamic stability.

South Africa offers a detailed historical context due to publicly accessible supply-scarcity indicators and geographically detailed load-shedding data [5–7]. The University of Pretoria dataset records hourly load-shedding stages in Tshwane from 2015 to 2023, including Hatfield, while Eskom provides data on demand, renewable generation, and unavailability [8,9]. These sources give scenario context but do not fully depict the electrical state or topology of a specific municipal feeder. NREL PVWatts offers modelled solar production profiles, not feeder measurements. This work uses these data to develop benchmark scenarios, maintaining these distinctions [10]. Physics-informed machine learning has been examined for power-system applications [11]. The present classifier uses power-flow-derived labels; its training objective does not enforce the network equations. Recent work addressed class imbalance in short-term voltage-stability assessment by generating synthetic minority-class samples with a conditional generative adversarial network and evaluated machine-learning algorithms for distribution-system voltage stability [12,13]. Building on these approaches, this study examines how operating scenarios, pre-simulation features, and cross-network reuse support allocation of a finite distribution-planning assessment budget. Here, the assessment budget is the number of candidate scenarios selected for detailed follow-up, as defined in Section 2.3; it is not the run time of the steady-state solver used for labelling. The study investigates the link between electrical risk identification and selecting operating states for detailed study, whether indicators can rank low-voltage cases, how rankings change with scenario descriptors, and performance retention when applying models across benchmarks. Precision–recall performance is emphasised because undervoltage cases form the minority class and the intended decision is how many of these cases can be retained within a fixed assessment budget. Figure 1 summarises the proposed application and its practical boundary. Historical data inform an offline scenario and training process. A deployed screening service would accept candidate operating conditions, check their compatibility with the trained model, rank them, and pass selected cases to a physical solver and an engineer. The benefit demonstrated here is concentration of low-voltage cases in the retained subset. The architecture is prospective; the experiments do not constitute deployment in a municipal distribution management system or a microgrid energy management system.

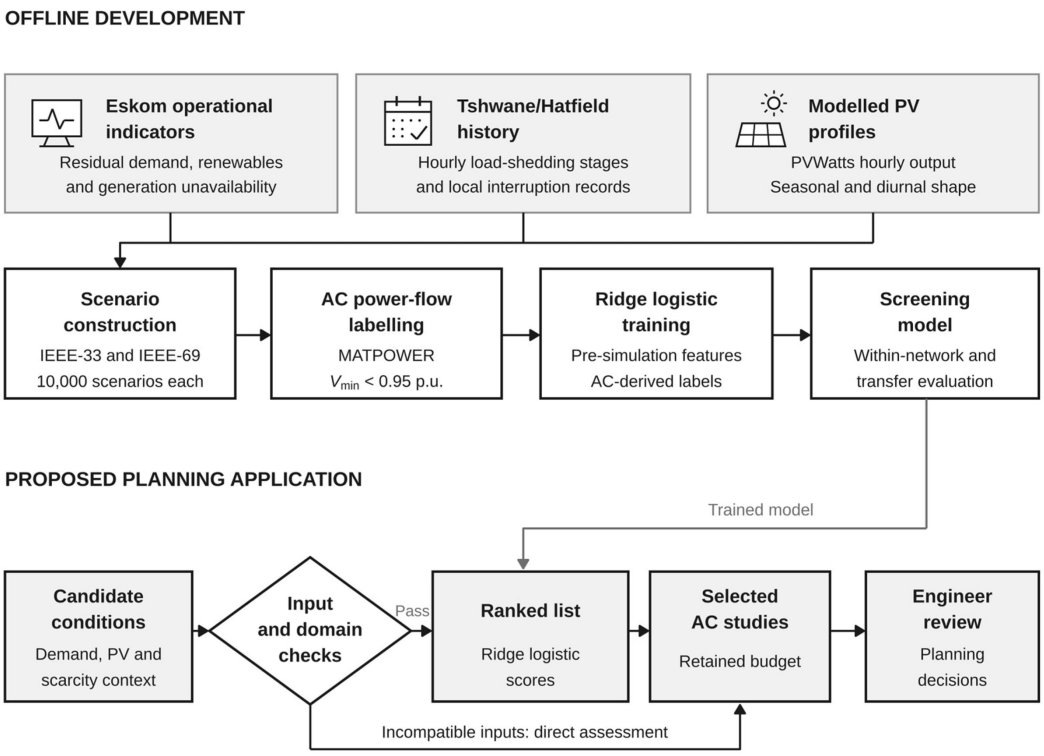



**Figure 1.** Study overview and prospective planning application.

The study makes three connected contributions to renewable-integrated distribution planning. First, it combines South African operational indicators, local Tshwane/Hatfield load-shedding histories and modelled PV profiles within a documented scenario-generation process, with voltage-risk labels obtained from AC power flow. Second, it distinguishes operational indicators from scenario-construction descriptors, allowing the information requirements and performance of alternative screening configurations to be evaluated explicitly. Third, it combines within-network assessments, IEEE-33-to-IEEE-69 transfers, and retained-budget analyses to determine how effectively low-voltage cases can be prioritised for detailed evaluation. Together, these contributions provide a digital engineering workflow linking urban operating context and renewable injections to a measurable allocation of network-assessment effort.

## 2. Background and Theoretical Framework

### *2.1. Low-Voltage Screening as a Defined Assessment Task*

The screening target is a voltage-magnitude condition at a solved operating point. A positive label means that at least one modelled bus is below the selected lower threshold. A negative label means only that this particular condition was not observed. It does not establish acceptable upper voltages, thermal loading, protection coordination, frequency response or dynamic recovery. In particular, the convergence of a numerical power-flow solver indicates that a solution was obtained under the solver settings; it is not proof of dynamic stability [3]. We let scenario $s$ produce voltage magnitudes $V_{i,s}$ at buses $i$ in a benchmark network. The threshold is fixed at 0.95 p.u. in the main experiment. It is a study criterion for prioritising low-voltage cases and not a claim of compliance with a universal statutory limit. A field application would need its own assessment criteria, voltage level,

measurement interval, and operating-condition definitions. For scenario $s$, the minimum bus-voltage magnitude over all modelled buses is given by

$$V_{\min,s} = \min_{i \in \mathcal{B}} V_{i,s}. \quad (1)$$

The scenario is then labelled by comparing this minimum with the screening threshold $\tau$, so that the binary undervoltage label is defined as

$$y_s = \mathbf{1}\{V_{\min,s} < \tau\}, \ \tau = 0.95 \text{ p.u.} \quad (2)$$

In Equations (1) and (2), $s$ identifies a scenario, $\mathcal{B}$ is the set of modelled bus identifiers, $i$ identifies a bus and $V_{i,s}$ is its solved voltage magnitude (p.u.). The indicator $\mathbf{1}\{\cdot\}$ equals one when its condition is true and zero otherwise; thus, $y_s = 1$ denotes undervoltage at the threshold $\tau$. Because $V_{i,s}$ is already a magnitude, no absolute-value operation is required in Equation (1), and the strict inequality in Equation (2) is used consistently throughout.

The indicator in Equation (2) is evaluated only for converged cases with finite voltages. A non-converged case remains unresolved and requires investigation; it must not be interpreted as a negative screening label. Sensitivity thresholds of 0.940, 0.955, and 0.960 p.u. examine how the chosen voltage criterion changes the labelled population. A tenth-percentile threshold is retained as a descriptive diagnostic. It is not used to impose a 10% low-voltage prevalence in the main experiment. An upper-voltage event is a separate target based on a scenario's maximum bus voltage exceeding an upper threshold. Raising the lower-voltage threshold in Equation (2) would still label low-voltage conditions and would not detect upper-voltage exceedances. The present negative class consequently means only that the lower-voltage screening criterion is not violated; it must not be interpreted as overall voltage compliance. No overvoltage prevalence or overvoltage-classifier performance is reported.

### *2.2. Supply Scarcity, Local Loading and Photovoltaic Injection*

Load shedding is a demand-reduction action associated with a supply shortfall. Its electrical effect depends on which loads are disconnected, what generation and reactive-power resources remain available, and how restoration is managed. Disconnecting demand can relieve voltage drop; load shedding cannot generally be assumed to cause undervoltage. National stage, municipal stage and local interruption status also describe different quantities. A national declaration does not imply that every feeder is interrupted during every hour [7,14].

In this study, load-shedding variables condition an assumed stress envelope for connected benchmark feeders. They do not operate breakers, de-energise buses or simulate islanding. Positive coefficients linking stage to local loading or load heterogeneity are scenario-design assumptions. Their role is to sample conditions associated with scarcity, not to identify a measured causal relationship between national stage and a particular bus voltage. This interpretation is necessary because the available public records do not contain the feeder-resolved measurements needed to identify that relationship directly.

Here, PV denotes photovoltaic generation. A photovoltaic connection location is distinct from a voltage-controlled power-flow PV bus, at which active power and voltage magnitude are prescribed. The added photovoltaic generation is represented as a prescribed active-power injection at a PQ bus, with no inverter reactive-power injection or absorption. It reduces the net active demand subject to the export cap, while the prescribed reactive demand is unchanged by the photovoltaic inverter. This representation is appropriate to the selected operating assumption and does not imply that photovoltaic plants are generally unable to regulate voltage. A plant operating in voltage-control mode would require a voltage setpoint and reactive-power capability limits; a Volt–VAR characteristic would

require its own control relation. These modes are not included in the present cases [15]. As a result, the study assesses the interactions of prescribed loads and a photovoltaic injection with the benchmark network; it does not quantify hosting capacity or the benefit of a particular inverter controller. The IEEE 33-bus feeder follows Baran and Wu [16], and the IEEE 69-bus feeder follows the radial system of Baran and Wu [17].

### *2.3. Physical Labels, Statistical Ranking and Retained-Case Budgets*

The framework couples a physical reference calculation to a statistical ranking model. The physical calculation determines the label for each training scenario. The classifier estimates a score from pre-simulation features. The score is used to order candidates, and the physical solver then checks the selected subset. The physics contribution is therefore in the labels and the network response. Unlike a physics-informed neural network, the ridge-logistic objective does not include differential-equation or power-balance residuals [18]. Figure 2 summarises the offline development sequence from data preparation and scenario construction to AC power-flow labelling and model training. This distinction matters: every development scenario must already have a reference label to evaluate pruning retrospectively. The reported reduction in retained cases therefore applies to a subsequent screening stage and not to the initial generation of the labelled training dataset.

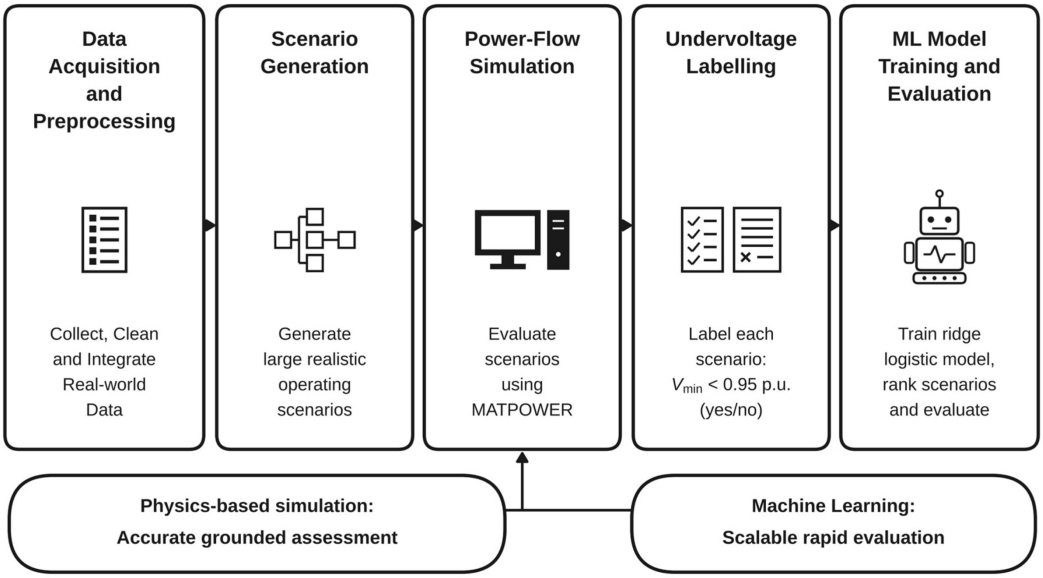



**Figure 2.** Theoretical framework for physics-labelled screening.

The assessment budget is the maximum number of candidate operating scenarios selected for detailed follow-up in a screening batch. In this study it is expressed as a case count, rather than a measured monetary cost or staff-time allowance. At the principal operating point, 3000 of 10,000 candidates are retained. Such a count could represent a limit for a downstream engineering review or in more expensive feeder studies, but these costs are not measured here. The demonstrated benefit is therefore the concentration of undervoltage cases in a prioritised subset. Every scenario in the present development libraries was already solved to obtain the reference labels; no saving is claimed for generating these libraries.

We let $N$ be the number of eligible candidate scenarios, $N_+$ the number with a positive label, $k$ the number retained for follow-up, and $TP_k$ the number of positive scenarios within the retained subset. Recall and precision distinguish the identity of many positives in a

small subset when recovering most positives in the whole candidate population. For a retained budget of $k$ scenarios, recall and precision are defined as

$$\text{Recall@}k = \frac{\text{TP}_k}{N_+}, \ \text{Precision@}k = \frac{\text{TP}_k}{k}. \tag{3}$$

The fractional reduction in the number of cases selected for follow-up, and the upper bound that this budget places on recall, are given by

$$R_{\text{calls}} = 1 - \frac{k}{N}, \ \text{Recall@}k \ \leq \ \min\left(1, \ \frac{k}{N_+}\right). \tag{4}$$

The denominators $N_+$ and $k$ in Equation (3) and $N$ and $N_+$ in Equation (4) are non-zero in all reported experiments. $R_{\text{calls}}$ is a dimensionless reduction in the number of cases selected for solver-based follow-up and not a measured reduction in elapsed time. For example, at $N_+ = 5107$ and $k = 3000$, the recall ceiling is $3000/5107 = 58.74\%$, as is consistent with Section 4.2.

Equation (4) links the retained budget to the positive-class prevalence. Capturing every positive case requires retaining at least as many scenarios as there are positives, even under perfect ordering. This necessary condition makes the engineering threshold and the retained-case budget jointly interpretable. The assessment therefore reports retained fraction, positive denominator, and missed cases alongside ranking performance.

## 3. Methodology

### *3.1. Data Sources, Provenance and Temporal Interpretation*

The data preparation combines national operational context, local historical interruption records, and a modelled solar profile. Hatfield is used as a geographical proxy for the University of Pretoria setting rather than as an electrical model of the campus. The underlying Tshwane source contains records for several areas [8]. The national operational variables describe the wider system and cannot substitute for measured active and reactive feeder demand [9]. Table 1 identifies the roles and interpretations of each component.

**Table 1.** Data components and their roles in the scenario study.

| Component | Information | Function |
| --- | --- | --- |
| Tshwane/Hatfield load-shedding records [8] | Hatfield load-shedding stage and active-event flag | Characterises local interruption conditions and supports construction of the outage-severity indicator. |
| National operational indicators retained in the integrated dataset [9] | Residual demand, renewable generation, generation-unavailability indicator $G$, national and municipal load-shedding stages, and manual load reduction | Provides national-system stress descriptors for scenario sampling and construction. |
| Derived operational indicators | Renewable-share indicator ($\rho_{\text{RE}}$) and outage-severity indicator $o$ | Provides normalised descriptors for stress-enriched sampling and machine-learning inputs. |
| PVWatts solar profile [19] | Normalised hourly PV output $v$, indexed by month and hour | Represents the seasonal and diurnal PV pattern used in scenario generation. |

The historical Tshwane workbook contains 77,088 rows, of which 39,864 hourly rows are retained in the integrated operational table. The latter covers 1 April 2019 to 17 October 2023, ending at 23:00. It contains 4393 rows flagged as active Hatfield load shedding and

local stage values from 0 to 6. The active-row share is therefore 4393/39,864 = 11.02%. This is a descriptive share of the processed record. It is different from the positive-label prevalence in the subsequently enriched simulation sample. The national operational columns used by the scenario generator were already present in the cleaned Tshwane table; the Eskom data [9] were used to check the consistency of these columns.

Figure 3 summarises the Hatfield load-shedding context retained for scenario calibration. The monthly stage-exposure map distinguishes the occurrence of different stages from the proportion of recorded hours associated with each stage. The accompanying calendar presents monthly mean severity over the same operational period. Together, these views provide a transparent temporal description of the local interruption indicators used in scenario construction.

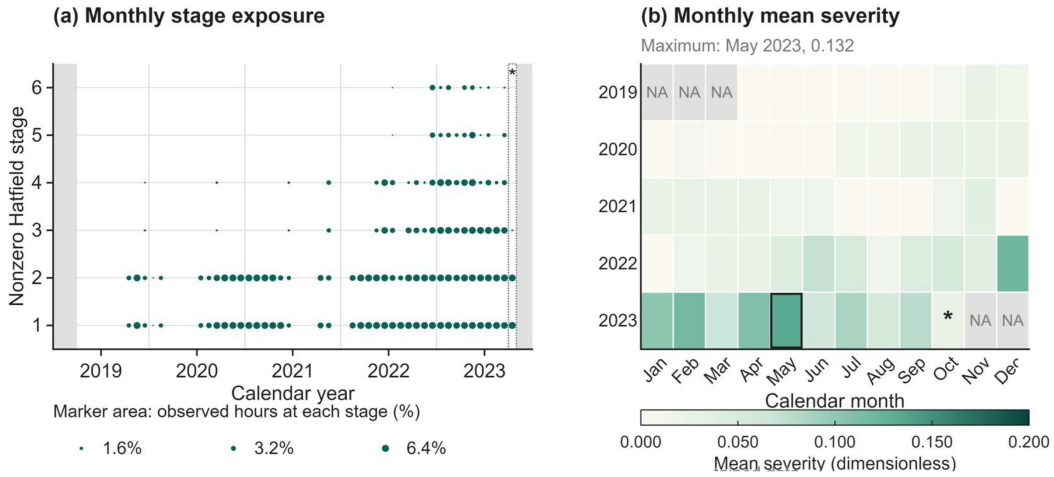



**Figure 3.** Integrated Hatfield load-shedding context used for scenario calibration. In (**b**), NA denotes months with no retained data and * denotes a partial calendar month (the record ends on 17 October 2023).

The two panels distinguish the frequency and level of recorded load-shedding exposure. Monthly mean severity can increase through a larger proportion of hours with active load shedding, higher stages during those hours, or both. The stage-exposure map makes these components visible, supporting interpretation of the aggregated severity measure. These historical summaries characterise the operating context supplied to the scenario generator; the resulting feeder-voltage response is subsequently determined through an AC power-flow assessment. The integrated operating record contains a generation-unavailability indicator, $G$, and a renewable-generation-to-residual-demand ratio, $\rho_{\mathrm{RE}} = 100\, P_{\mathrm{RE}}/P_{\mathrm{res}}$ (%), where $P_{\mathrm{RE}}$ is total renewable generation and $P_{\mathrm{res}}$ is residual demand, both in MW. Their percentile-normalised forms provide contextual descriptors for scenario construction. $\rho_{\mathrm{RE}}$ is a ratio of national operational quantities; the local PV fraction is specified separately when constructing feeder injections.

PVWatts supplies an hourly solar-output shape, with records grouped by calendar month and hour [19]. A scenario samples a PV record from the corresponding month–hour group of its selected operational record. This preserves a diurnal and seasonal association without asserting that the PV output was observed at the same historical timestamp. If the requested month–hour group is empty, the implementation samples from the complete modelled PV profile. PVWatts supplies the modelled production conditional on its weather input and photovoltaic-system assumptions [20]. In this study, variation in the supplied hourly profile is sampled within month–hour groups, and the scenario generator also perturbs the local photovoltaic fraction using the noise terms documented in Table 2 and Appendix A. This introduces variability among the constructed operating points. It is not a

complete stochastic irradiance model: independent scenario sampling does not reproduce a chronological cloud-ramp process, temporal autocorrelation, spatially correlated irradiance, or forecast-error distributions. The present analysis therefore addresses variability represented by the sampled hourly profile and the stated scenario perturbations, rather than validated sub-hourly photovoltaic uncertainty. Figure 4 combines the modelled solar-production shape with the dependence structure of the integrated operational variables. The correlation matrix provides a descriptive check on dependence among these inputs and is not a validation of the benchmark load-mapping coefficients. As shown in Figure 4b, Hatfield stage and the derived outage-severity indicator have a displayed Pearson correlation of 1.00. Total renewable generation and the renewable-share proxy are also strongly related, with a displayed correlation of 0.98. These dependencies motivate the subsequent feature-set comparison and constrain interpretation of individual feature importance.

**Table 2.** Scenario-generation settings used in the IEEE-33 and IEEE-69 case studies.

| **Setting** | **IEEE-33** | **IEEE-69** |
|---|---|---|
| Number of scenarios | 10,000 | 10,000 |
| Random seed | 42 | 69 |
| Benchmark network | MATPOWER case33bw | IEEE 69-bus radial distribution test feeder |
| Low-voltage criterion | $V_{\min} < 0.95$ p.u. | $V_{\min} < 0.95$ p.u. |
| Load-multiplier bounds, $L$ | 0.28–1.08 | 0.26–1.03 |
| Reactive-load-multiplier bounds, $Q$ | 0.76–1.04 | 0.74–1.00 |
| Noise standard deviations, ($\sigma L$, $\sigma Q$, $\sigma f$) | (0.020, 0.015, 0.040) | (0.018, 0.013, 0.040) |
| Scenario-sampling rule | Equation (6) | Equation (6) |
| Scenario-envelope definitions | Equations (A1), (A3) and (A5) | Equations (A2), (A4) and (A5) |

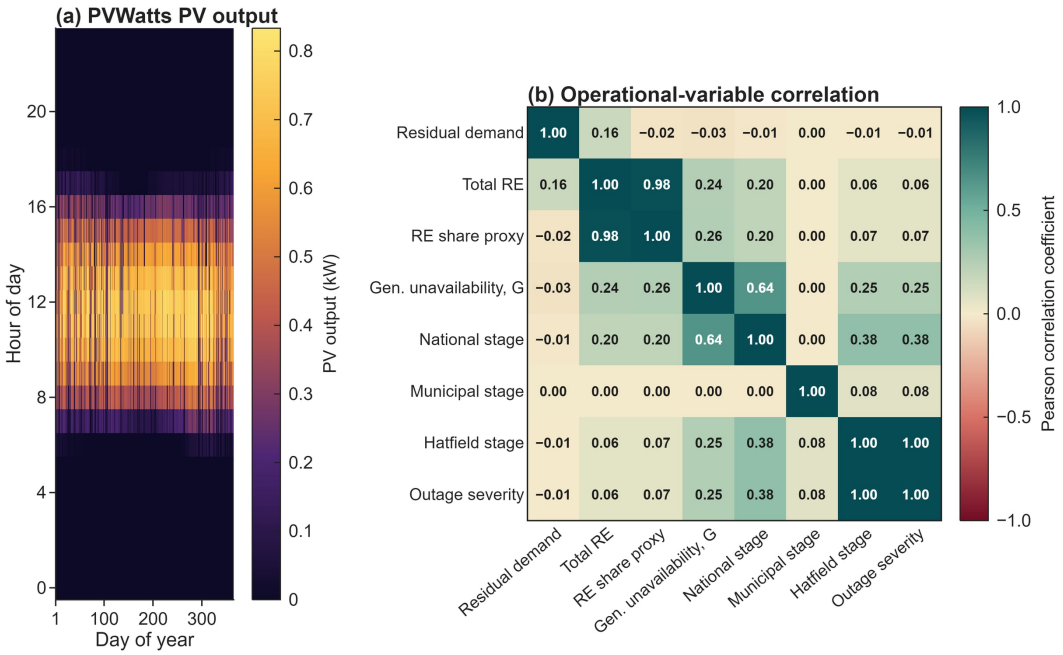



**Figure 4.** Solar profile and dependence among operational indicators.

### 3.2. Stress-Enriched Sampling and Scenario Construction

National residual demand, $G$, the renewable-share indicator, and manual load reduction are normalised using their fifth and ninety-fifth percentiles in the operational calibration table. For a source variable with value $x$, the percentile-normalised value, clipped to the interval $[0, 1]$, is given by

$$z_x = \text{clip}\left(\frac{x - q_{0.05,x}}{q_{0.95,x} - q_{0.05,x}}, 0, 1\right). \tag{5}$$

where $q_{0.05,x}$ and $q_{0.95,x}$ are the fifth and ninety-fifth percentiles of the source variable in the historical calibration table, and $\text{clip}(u, a, b) = \min(\max(u, a), b)$. All resulting indicators are dimensionless. The transformation requires a non-zero percentile range, $q_{0.95,x} > q_{0.05,x}$. Hatfield, national, and municipal stages are divided by eight and clipped to $[0, 1]$. Dividing by eight is a numerical normalisation; it does not imply that stages 7 or 8 occurred in the retained Hatfield record. The normalised PV shape $v$ is the sampled output divided by the maximum of the modelled PV profile.

We let $r$ and $g$ denote the normalised residual demand and generation unavailability; $h$, $n$, and $m$ denote the normalised Hatfield, national, and municipal stages; $a$ indicates active local load shedding; and $o$ is the derived outage-severity indicator. This processing defines $o = ah$; thus, $H$ denotes the unnormalised Hatfield stage. Historical rows are sampled using replacement with the weights in Equation (6), normalised by their sum across eligible records. Weighting gives scarce operating conditions greater representation in the training scenario space. No likelihood correction is subsequently used to reinterpret the sample as the natural frequency distribution of feeder events. Each eligible historical hourly record $t$ is assigned the stress-enriched sampling weight

$$w_t = 1 + 0.65r_t + 0.55g_t + 1.10o_t + 0.70\mathbf{1}\{H_t \geq 4\} + 0.35a_t. \tag{6}$$

In Equation (6), subscript $t$ denotes the value in record $t$; $r_t$ and $g_t$ are the percentile-normalised residual demand and generation unavailability; $H_t$ is the unnormalised Hatfield stage; $a_t$ is the active-event flag; $h_t = H_t/8$, representing the normalised stage; and $o_t = a_t h_t$. The probability of drawing record $t$ is $w_t/\Sigma_u w_u$, where the sum runs over all eligible records. The constants are assumed to be sampling-design coefficients and not calibrated causal effects.

Each selected record produces a load multiplier $L$, a reactive-load modifier $Q$, a local PV fraction $f$, and a composite stress descriptor $S$. Gaussian perturbations introduce scenario-to-scenario variability. The final envelopes differ slightly between the two networks. Table 2 lists the final settings, and Appendix A gives the scenario-envelope expressions. These are assumed benchmark mappings, not parameters fitted against measured Hatfield feeder voltages.

All multipliers and noise standard deviations are dimensionless. Pilot simulations informed the feeder-specific scenario envelopes, which were then fixed for the reported 10,000-scenario evaluations. The resulting positive-class prevalences reflect both the stress-enriched sampling procedure and these modelling choices. The final coefficients and bounds specify the operating populations to which the reported screening results apply.

Photovoltaic generation is allocated to buses 6, 13, 18, 25, 30, and 33 in IEEE-33, and buses 11, 18, 27, 35, 46, 52, 61, and 69 in IEEE-69. Allocation weights are proportional to the positive base active demands at the selected buses. Requested aggregate PV power is the product of yjr total base active demand, the scenario local PV fraction, and the normalised PV shape. A lower bound on the net active demand limits export at each photovoltaic connection bus to 0.25 times that of the total base active demand divided by the number of photovoltaic connection buses. The descriptor $P_{\text{PV,requested}}$ (Equation (A6))

denotes requested aggregate photovoltaic power; the injection applied to the network is determined by the bus-allocation and export-cap calculations and can be smaller after clipping. Figure 5 shows the benchmark feeder structures and the buses selected for local PV injection. Energised branches are distinguished from branches stored as out of service in the MATPOWER network representation, including normally open ties where defined. The reference bus and photovoltaic connection buses are highlighted explicitly.

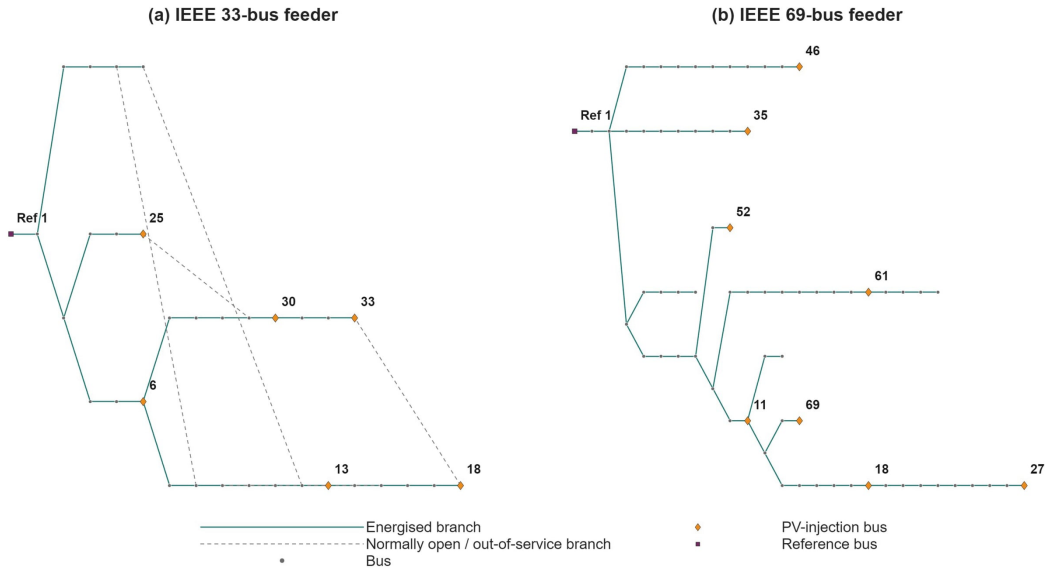



**Figure 5.** Benchmark feeder topologies and photovoltaic connection locations.

### *3.3. Network Modification and Physical Labels*

The base active and reactive demands are scaled independently of the public system-demand units. National MW values first enter the dimensionless mapping; they are not injected directly into a small distribution benchmark. Before the export cap is applied, the net active demand and the reactive demand at bus $i$ in scenario $s$ are given by

$$P_{i,s}^{\text{net}} = L_s \eta_{i,s} P_i^0 - A_{i,s}, \; Q_{i,s}^{\text{d}} = L_s Q_s \eta_{i,s} Q_i^0. \tag{7}$$

In Equation (7), ${P_i}^0$ (MW) and ${Q_i}^0$ (MVAr) are the base active and reactive demands at bus $i$, $L_s$ is the dimensionless active-demand multiplier, $Q_s$ is a dimensionless additional reactive-demand modifier rather than a reactive power, $\eta_{i,s}$ is the dimensionless local heterogeneity factor (Appendix A), and $A_{i,s}$ (MW) is the allocated photovoltaic active power before the export cap is applied. Positive ${P_{i,s}}^{\text{net}}$ denotes consumption. The common factor $L_s$ scales both active and reactive demand. The export cap (Section 3.2) then limits the applied net active demand, so the applied photovoltaic injection can be smaller than the requested allocation. At a photovoltaic connection bus with no conventional generator, the active and reactive injections used in Equation (8a,b) are the negatives of the applied net active demand and of the prescribed reactive demand, each divided by the case MVA base; photovoltaic power is therefore not counted twice.

Equation (7) specifies the photovoltaic contribution as active power only. Adding a photovoltaic injection does not change the connection bus into a voltage-controlled PV bus. If a photovoltaic inverter regulates voltage, its reactive output becomes part of the operating-point solution, subject to the device's reactive-power limits. This can change minimum voltages, the location of the minimum, the binary labels and the resulting scenario ranking. The reported classifier performance therefore applies to the adopted prescribed-injection model and has not been validated for voltage-regulating photovoltaic plants. No quantitative comparison between these operating modes is claimed.

Heterogeneity includes a Gaussian term and stage- and unavailability-dependent terms multiplied by the normalised bus index. Bus index is an assumed spatial weighting, not electrical distance. The branch resistances, reactances, reference-source settings and topology are inherited unchanged from the selected network case. No outage-specific branch switching or closed-loop voltage control is applied. These assumptions provide a controlled learning problem but restrict the interpretation of the resulting spatial voltage patterns.

Bounding the reference-bus voltage does not generally bound every PQ-bus voltage in a feeder containing distributed active-power injections. The usual voltage-drop argument assumes predominantly downstream consumption; a photovoltaic connection with a negative net active demand can produce reverse active-power flow and local voltage rise [21]. The implemented export cap limits the allowed magnitude of injection but does not impose a voltage constraint. Upper-voltage behaviour therefore cannot be excluded from the present assumptions solely by constraining the reference source. Its prevalence in the final scenario libraries has not been quantified in this study.

AC power flow is solved in MATLAB using MATPOWER [22]. Each solved operating point satisfies the AC nodal power-balance equations [15,22]. For bus $i$ in scenario $s$, the net active-power injection satisfies

$$p_{i,s} = \sum_{j \in \mathcal{B}} V_{i,s} V_{j,s} \left[ G_{ij} \cos\left(\delta_{i,s} - \delta_{j,s}\right) + B_{ij} \sin\left(\delta_{i,s} - \delta_{j,s}\right) \right]. \tag{8a}$$

Likewise, the net reactive-power injection satisfies

$$q_{i,s} = \sum_{j \in \mathcal{B}} V_{i,s} V_{j,s} \left[ G_{ij} \sin\left(\delta_{i,s} - \delta_{j,s}\right) - B_{ij} \cos\left(\delta_{i,s} - \delta_{j,s}\right) \right]. \tag{8b}$$

In Equation (8a,b), $p_{i,s}$ and $q_{i,s}$ are the net active and reactive injections at bus $i$ in scenario $s$ (p.u.), $\delta_{i,s}$ is the voltage angle (rad), $G_{ij}$ and $B_{ij}$ are the real and imaginary parts of the bus-admittance matrix entries (p.u.), and $j$ runs over all buses in $\mathcal{B}$. Injections in MW and MVAr are divided by the case MVA base, such that MW quantities are not mixed with the per-unit admittance matrix. The source indicator $G$ of Section 3.1 is distinct from the admittance entries $G_{ij}$.

Each case is solved with the MATPOWER runpf function using the default Newton–Raphson settings of MATPOWER, with the console output suppressed. For each case, the stored outputs comprise the convergence status; per-scenario solver times; the minimum, mean and maximum voltages; the bus with the minimum voltage; and the full bus-voltage arrays. The main label then follows Equation (2).

### *3.4. Feature Sets and Classifier Configuration*

Four feature sets distinguish operational information from variables created inside the scenario generator. Operational minimal (OM) contains the residual demand, $\rho_{\mathrm{RE}}$, $G$, national stage, Hatfield stage, outage severity, and PV shape. These seven indicators are more readily associated with operational data services, but the name does not imply a mathematically minimal or independent representation. Under the implemented construction, outage severity equals the normalised Hatfield stage: a positive stage activates the event flag, whereas a zero stage makes their product zero. These columns therefore represent dependent information within the operational feature set. Operational without derived outage severity (ON) contains the residual demand, $\rho_{\mathrm{RE}}$, $G$, national stage, municipal stage, Hatfield stage, and PV shape. Relative to OM, it both removes outage severity and adds municipal stage. It is therefore an alternative representation rather than a controlled single-feature ablation. Scenario descriptors only (SDO) contains $L$, $Q$, requested photovoltaic power, local photovoltaic fraction and $S$. Hybrid reduced (HR) contains the normalised residual demand, normalised $G$, normalised renewable share, normalised Hatfield stage,

PV shape, $L$, $Q$, and local photovoltaic fraction. The latter two configurations are intended for a scenario engine that already knows its construction variables.

None of the feature lists includes solved bus voltages, the worst-bus output, convergence-derived voltage information, or another risk-label column. However, excluding post-simulation variables does not remove the dependence between the classifier inputs and the rules used to construct labels. Descriptor-based performance is consequently interpreted as learning the selected scenario-to-voltage relationship. It is not a theoretical upper bound on every possible model and does not show that the same predictors are sufficient for an unobserved physical feeder.

The principal model is ridge-logistic regression, as implemented in MATLAB 2026a using fitclinear [23] with the logistic learner, ridge regularisation parameter $10^{-4}$, the L-BFGS solver and an iteration limit of 1000. Ridge regularisation adds a penalty on the squared coefficient magnitudes to the logistic loss and shrinks the coefficient estimates towards zero [24]. The regularisation parameter is fixed at the stated value for the reported experiments. All reported feature-set, model-reuse, and retained-budget results use this ridge-logistic configuration.

### *3.5. Evaluation Protocol and Dependence Limits*

Within each network, classification performance is evaluated using 20 stratified hold-outs, with approximately 70% of eligible scenarios assigned to training and 30% to testing. Predictor means and standard deviations are estimated from the training partition and applied to the test partition. Records with unavailable or non-finite inputs are excluded, and effectively constant predictors are screened before partitioning.

Class weights $w_c = N/(2N_c)$ are calculated from the full eligible labelled dataset, where $N_c$ is the number of scenarios in class $c \in \{0, 1\}$, and are assigned to the training observations. Different feature sets use different random-seed sequences; their reported mean differences are therefore descriptive comparisons rather than paired statistical tests.

Because the class-frequency weights were computed using the complete eligible labelled dataset, held-out labels contributed to the class-frequency estimate used for weighting. The resulting scores are interpreted as a retrospective benchmark of performance rather than a strictly isolated prospective validation. This limitation is additional to the dependence between scenarios sharing historical source records.

Pruning within a network uses five-fold out-of-fold scores; the model producing a score for a row was trained without that row. These pooled scores support retrospective prioritisation of the full scenario population. They are not predictions from one independently deployed model on a new 10,000-case batch. For transfer, a model is fitted using IEEE-33 data and applied to IEEE-69 features standardised with the IEEE-33 training statistics. The corresponding classifier has not been fitted to IEEE-69 labels, but the IEEE-69 scenario envelope was adjusted during pilot design. Both datasets also draw from the same operational record.

The classification metrics are area under the receiver operating characteristic curve (ROC-AUC), area under the precision–recall curve (PR-AUC), recall, and associated confusion-matrix measures. Recall is TP/(TP + FN), and false-negative rate is one minus recall. The reported ROC-AUC and PR-AUC values are calculated by trapezoidal integration of score-ordered curves. PR-AUC therefore denotes the area obtained using this integration convention and is distinct from average precision [25].

Random scenario splits quantify interpolation within the constructed dataset. They do not establish independence between historical operating periods, because repeated scenarios can share an operational timestamp or a PV source record. A temporally blocked or source-record-grouped evaluation has not been demonstrated. This limits both future-

period and cross-network conclusions, even where the reported numerical performance is high.

### *3.6. Pruning, Baselines and Interpretation*

Candidate scenarios are ordered by descending ridge-logistic score. The retained fraction is varied between 1% and 50% of the candidate library, and the 30% budget is used for the principal quantitative comparison. Recall@$k$ and Precision@$k$ follow Equation (3). Every deferred positive case is counted as missed instead of being removed from the denominator. The cases outside the retained subset are deferred by the ranking and are not certified as voltage compliant. In particular, the operational rankings at the principal budget leave 89 IEEE-33 undervoltage cases, 72 IEEE-69 cases and 99 transfer cases outside the selected subset. These missed counts limit the use of the ranking as an exhaustive safety screen. At a fixed budget, a random ranking has expected recall $k/N$, which provides a transparent reference without requiring an additional learned model. The heuristic comparison reported in Section 4.4 examines simple rankings using load scaling, the composite stress descriptor, generation-unavailability stress, stage, outage severity, and a combined engineering score at a 10% budget. This heuristic diagnostic is reported for IEEE-33 and at the 10% budget only. It illustrates the behaviour of simple engineering rankings at a stringent budget and is not a matched comparison with the learned rankings at the principal 30% budget or on IEEE-69. These are inexpensive heuristic baselines. They are not evaluations of a fast voltage stability index, an L-index, or continuation-power-flow margin. A ranking based on solved minimum voltage would instead be an oracle requiring the very AC evaluations that screening aims to reduce. The manuscript therefore does not claim superiority to classical voltage-stability methods.

Permutation importance is calculated by permuting one test-set feature and recording the resulting decrease in PR-AUC. It describes how much the fitted predictor relies on that feature under the adopted experiment. It does not estimate a causal effect on a real feeder. Strong predictor dependence can redistribute apparent importance among correlated variables, a limitation documented in the variable-importance literature [26]. The importance analysis uses one holdout partition and one permutation per feature; it does not provide uncertainty intervals across repeated permutations.

### *3.7. Use of Generative Artificial Intelligence*

Generative artificial intelligence assistance was used when checking the results and during revision to support drafting and preparation of mathematical notations. The tools are identified in the Acknowledgments.

## 4. Results and Discussion

### *4.1. Voltage Distributions and Network Response*

The final runs contain 10,000 converged power flows for each feeder. Minimum bus voltages span 0.9083–0.9986 p.u. for IEEE-33 and 0.9083–0.9962 p.u. for IEEE-69. At the main 0.95 p.u. threshold, 2600 IEEE-33 scenarios and 2282 IEEE-69 scenarios are positive. These correspond to sample prevalences of 26.00% and 22.82%. Figure 6 shows the two distributions and the common screening threshold, enabling the classifier results to be interpreted against the underlying label populations.

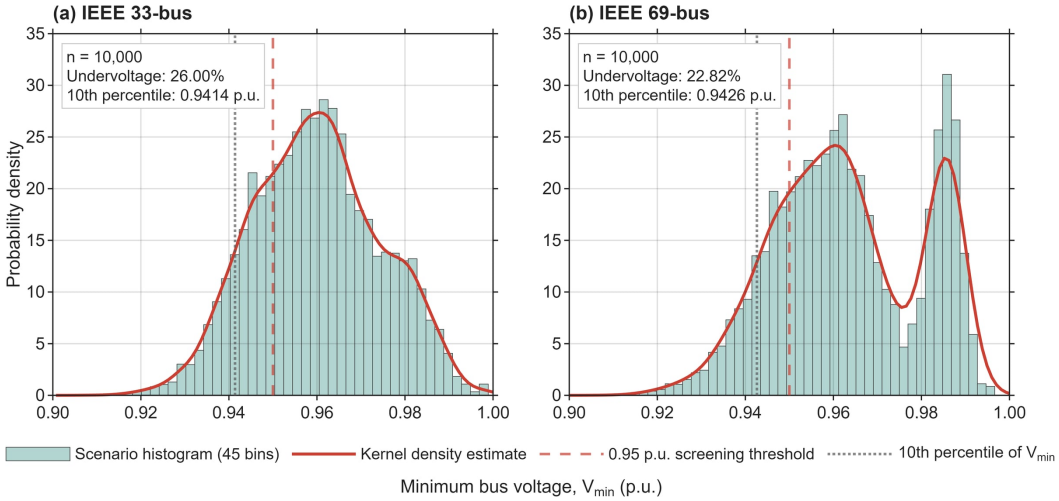



**Figure 6.** Distributions of scenario-level minimum bus voltages for 10,000 scenarios in each feeder.

The voltage distributions characterise the constructed operating populations, whose differences reflect both the feeder models and the adopted demand and photovoltaic mappings. Figure 7 illustrates three actual scenarios from each library: the scenario with the lowest minimum voltage, an actual scenario nearest to the median of the scenario minima, and the scenario with the highest minimum voltage. These rules select scenarios; they do not select a predetermined weak bus. The bus attaining the minimum can change with the spatial pattern of active and reactive demand and photovoltaic injection. Bus number is a node identifier and not cumulative electrical distance, so consecutive plotted bus numbers can belong to different laterals. In the standard 33-bus topology, for example, bus 19 connects to bus 2 rather than to bus 18 [16].

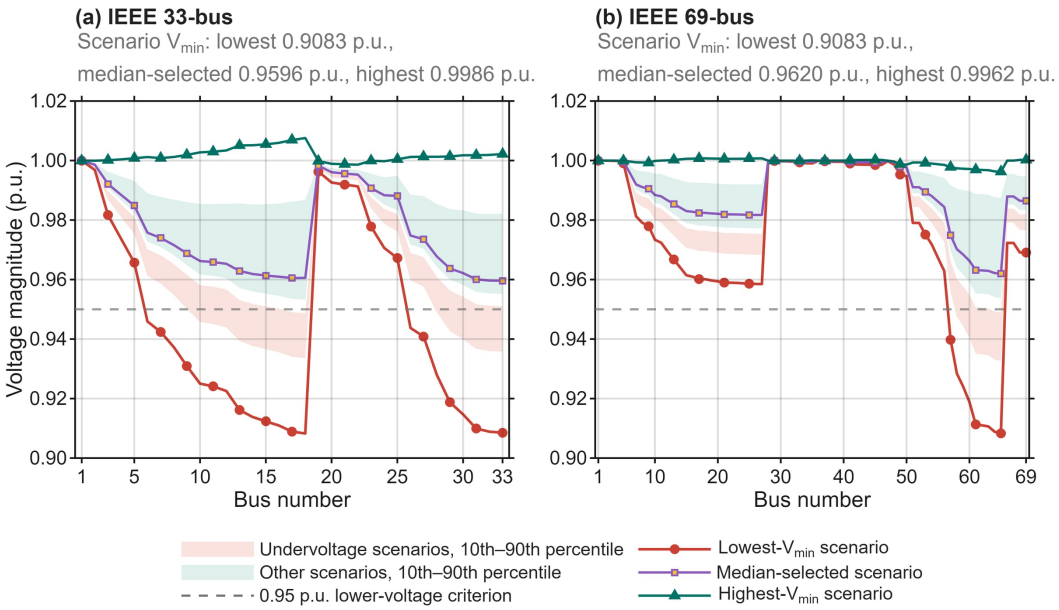



**Figure 7.** Representative bus-voltage profiles for the IEEE-33 and IEEE-69 scenario libraries.

The voltage profiles provide the spatial information underlying the scenario-level risk label. Minimum voltage identifies the most exposed bus, whereas the complete profile shows how the voltage reduction is distributed across the network. This distinction matters because an acceptable feeder-average voltage can coexist with a local undervoltage condition. Selecting representative scenarios according to their minimum voltage therefore connects the binary screening outcome to an identifiable electrical operating state.

The profile differences are consistent with the combined active- and reactive-demand representation in Equation (7). PV injection reduces net active demand at the selected buses, while reactive demand is prescribed separately. Consequently, the voltage response depends on the distribution of demand and PV injection together with the feeder impedances and local loading multipliers. Aggregate renewable availability alone does not specify this spatial response.

In the IEEE-33 panel of Figure 7, the lowest voltage of the minimum-scenario profile occurs at bus 18, the terminal bus of the main feeder (buses 1–18), which has the largest cumulative series impedance from the reference bus. The highest voltages of the maximum-scenario profile occur along the same feeder section, where several bus voltages exceed the 1.00 p.u. reference-bus voltage. Both extremes therefore arise towards the remote end of the main feeder: the same cumulative impedance produces a voltage drop when net demand flows towards the feeder end and a voltage rise when photovoltaic injection at connection buses, such as buses 13 and 18, exceeds local demand and reverses the flow [21]. These profiles also show that, in the present model, the net photovoltaic export can raise PQ-bus voltages above the reference voltage, as discussed in Section 3.3.

For a prospective urban distribution-planning study, these profiles could provide a second stage of assessment after scenario ranking. The screening model would identify operating conditions requiring attention, and the corresponding AC power-flow solution would identify the buses at which detailed investigation is warranted. Engineers could then examine local voltage support, demand flexibility or storage operation using the retained scenarios as study inputs. This use has not been demonstrated on a municipal feeder.

Figure 8 relates the solved minimum voltage to load scaling, normalised generation-unavailability stress, and normalised Hatfield stage for the two benchmark feeders. Each point represents a converged scenario, and its colour indicates the minimum voltage obtained from AC power flow. The visualisation connects the historical stress descriptors to the electrical response of the generated operating states. The association between increased load scaling and reduced minimum voltage is consistent with the higher network demand imposed by the scenario model. However, the three displayed coordinates do not completely determine a scenario, as PV injection, reactive-demand scaling, and bus-level heterogeneity also vary. Differences in voltage among scenarios occupying similar regions of the displayed space can therefore arise from operating variables outside those three axes. This motivates a multivariable screening representation.

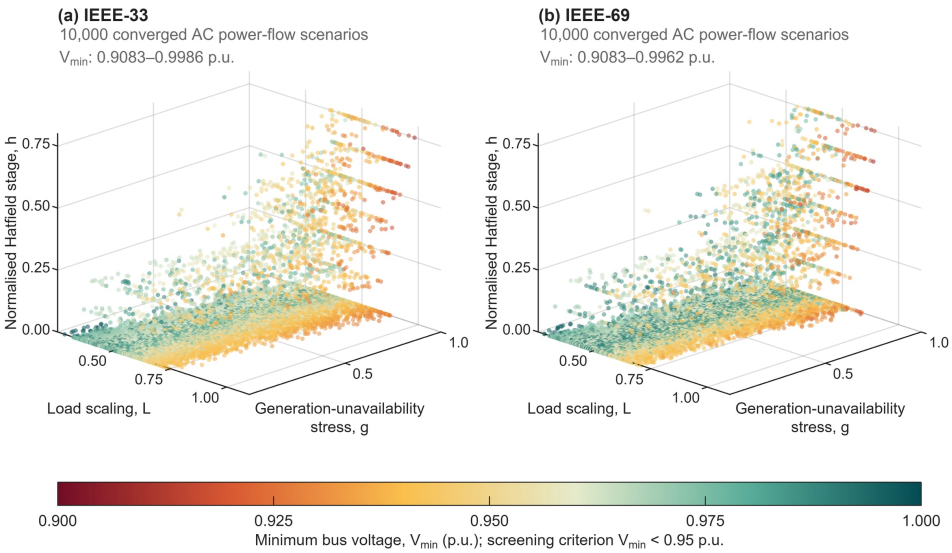



**Figure 8.** Operating-space maps of simulated minimum voltage.

The map also distinguishes the roles of operational context and physical assessment. Generation unavailability and load-shedding stage condition the scenario-generation process, while the voltage-risk label is obtained from the resulting AC solution. Their relationship with voltage is consequently interpreted within the specified scenario mapping. The framework provides a traceable connection from public operational indicators, through prescribed demand and renewable injections, to an electrical assessment criterion.

For prospective planning studies, this connection makes the operating assumptions behind a risk ranking inspectable. A planner can see which demand, PV availability, and supply-scarcity conditions populate the low-voltage region and use these to define feeder studies. The feature-set comparison then assesses how well the pre-simulation information captures this relationship.

### *4.2. Threshold Sensitivity and Within-Network Screening*

The voltage threshold materially changes the positive population. For IEEE-33, the reported counts are 822 at 0.940 p.u., 2600 at 0.950 p.u., 3758 at 0.955 p.u., and 5107 at 0.960 p.u. The corresponding shares are 8.22%, 26.00%, 37.58%, and 51.07%. The tenth-percentile diagnostic is 0.941393 p.u. Figure 9 combines threshold-dependent changes in the labelled population with the within-network comparison of operational and scenario-descriptor features.

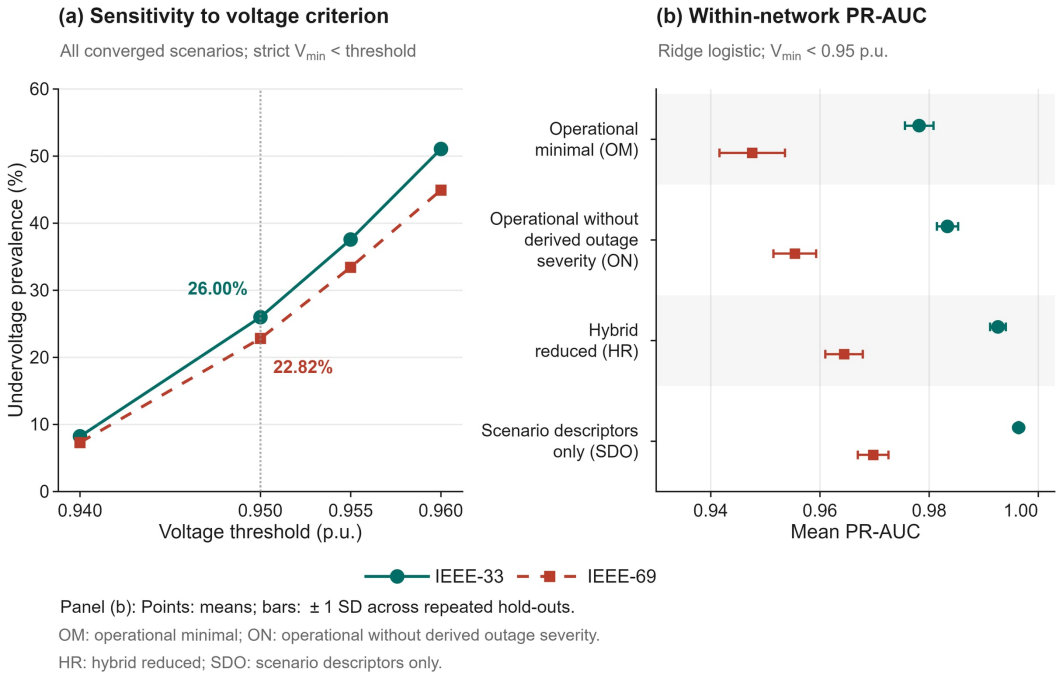



**Figure 9.** Label sensitivity and within-network feature-set performance.

These values show that label prevalence depends strongly on the chosen screening criterion; they do not demonstrate classifier robustness to alternative labels unless the learning evaluation is also repeated for each criterion. The threshold analysis examines which operating states enter the undervoltage class, while the classification comparison evaluates screening performance at the main criterion of 0.95 p.u. Together, these analyses distinguish the effect of the engineering assessment criterion from the effect of predictor information. The IEEE-33 threshold results show that a relatively small change in the screening criterion can substantially change the number of scenarios requiring attention. Increasing the threshold from 0.950 to 0.955 p.u. adds 1158 positive cases, corresponding to 11.58% of the scenario library. These cases occupy the voltage interval immediately

above the main threshold and illustrate how a more precautionary criterion expands the assessment workload.

Threshold selection and the retained-scenario budget are therefore coupled planning decisions. For example, at 0.960 p.u., 5107 of the 10,000 IEEE-33 scenarios are positive. A budget of 3000 cases could then capture at most 58.74% of the positives, even with perfect ranking, as follows directly from Equation (4). This is a calculated budget bound, independent of classifier performance. It demonstrates why a retained fraction must be interpreted together with the selected voltage criterion.

At the main 0.95 p.u. criterion, the feature-set comparison examines how much screening information is available before running the power flow. Operational features represent the historical operating context and PV availability, whereas scenario descriptors specify additional quantities used to construct the electrical operating point. Table 3 summarises the corresponding within-network performances.

**Table 3.** Reported within-network ridge-logistic performance at $V_{min} < 0.95$ p.u.

| Network | Feature Set | Mean PR-AUC | Mean Recall | FNR |
|---|---|---|---|---|
| IEEE-33 | Operational minimal | 0.9782 | 0.9606 | 0.0394 |
| IEEE-33 | Scenario descriptors only | 0.9964 | 0.9873 | 0.0127 |
| IEEE-69 | Operational minimal | 0.9476 | 0.9420 | 0.0580 |
| IEEE-69 | Scenario descriptors only | 0.9697 | 0.9623 | 0.0377 |

Means refer to 20 stratified 70/30 holdouts. FNR is calculated as one—mean recall. These are classification metrics and do not indicate top-budget recall.

Both representations achieve strong within-network discrimination, with scenario descriptors providing higher mean PR-AUC in each feeder. Relative to OM, the descriptor representation increases PR-AUC by 0.0182 on IEEE-33 and 0.0221 on IEEE-69. The corresponding increases in classification recall are 2.67 and 2.03 percentage points. These differences quantify the additional predictive information available when load scaling, reactive-demand scaling and requested PV injection are known during scenario construction.

The comparison distinguishes two screening configurations. OM evaluates how effectively contextual operational indicators can prioritise the generated operating states. SDO evaluates the additional value of variables that the scenario generator itself uses to construct each operating point: the load and reactive-demand multipliers, the requested photovoltaic power, the local photovoltaic fraction, and the composite stress descriptor. These are pre-simulation inputs and not solved voltages or labels (Section 3.4). However, because they define the inputs to the physical calculation directly, the SDO scores are conditional on the scenario generator, the stress-enriched sampling, and the prescribed benchmark network models. The SDO results therefore cannot be read as screening indicators on real feeders, as such construction variables would be uncertain or unavailable. The OM results are likewise retrospective benchmark results; neither representation has been validated with field data.

### *4.3. Model Reuse from IEEE-33 to IEEE-69*

The model-reuse experiment tests a different question from repeated validation: whether a classifier trained on IEEE-33 retains useful ranking when directly applied to IEEE-69. Figure 10 compares transfer PR-AUC with the IEEE-69 within-network reference for each feature set. OM declines from 0.9476 within IEEE-69 to 0.9376 under transfer, an absolute difference of 0.0099. SDO declines from 0.9697 to 0.9596, a difference of 0.0102. These are descriptive differences between distinct training and evaluation protocols. Both differences are computed from unrounded values, as shown in Figure 10.

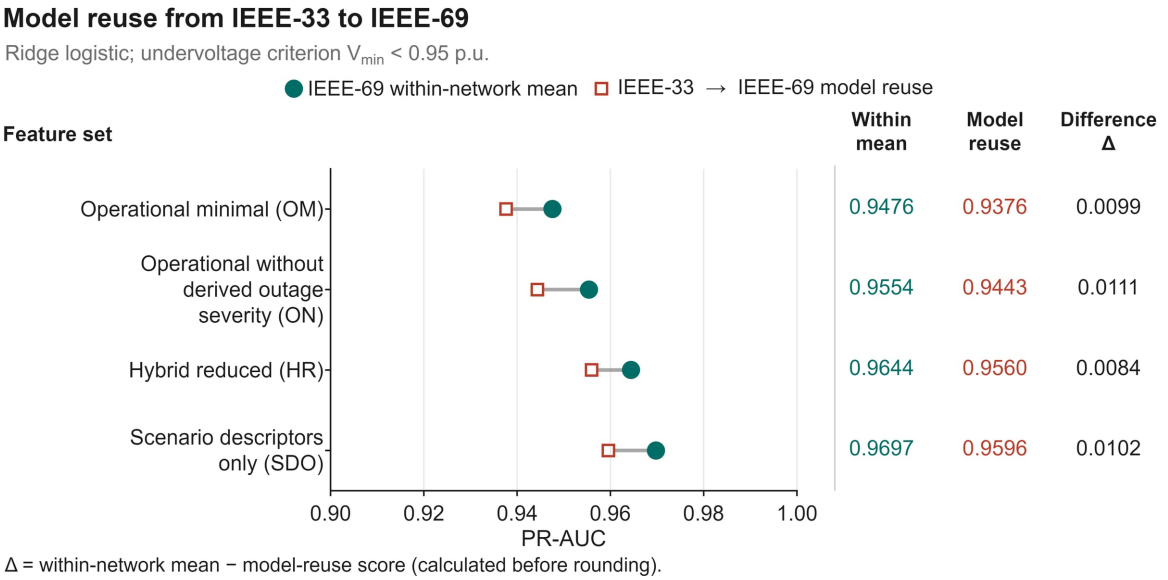



**Figure 10.** Model reuse from IEEE-33 to IEEE-69 within the shared scenario-modelling framework. The IEEE-33 classifier and standardisation parameters are applied to IEEE-69 without refitting; both libraries use the same historical operational source and related scenario-generation rules.

These results show that the useful separation of low-voltage and non-low-voltage scenarios is retained when the IEEE-33 classifier and its standardisation parameters are applied to IEEE-69 without fitting the classifier to IEEE-69 labels. The experiment is therefore a useful test of model reuse between two network models. It is not a test of generalisation to an unseen operating distribution. Both libraries draw on the same historical operational record and related scenario-generation rules, and the IEEE-69 scenario envelope was adjusted during pilot design. The results therefore do not establish generalisation to an independently sampled operating distribution, a future period, or a real feeder.

The operational and descriptor representations behave differently when their scores are converted into binary classifications. OM combines transfer PR-AUC of 0.9376 with recall of 0.9553. SDO achieves the higher PR-AUC of 0.9596, but its classification recall is 0.8002. Thus, a representation can preserve strong overall discrimination while placing fewer positive cases above the transferred decision threshold. This distinction is relevant when choosing between a binary alarm and a ranking used to allocate a simulation budget.

Table 4 reports both measures for all four representations. ON provides the highest classification recall among the configurations shown, while SDO provides the highest PR-AUC. These outcomes support selecting the representation according to the intended screening decision and available inputs.

**Table 4.** Reported direct IEEE-33-to-IEEE-69 ridge-logistic model-reuse performance.

| Feature Set | PR-AUC | Recall | FNR |
|---|---|---|---|
| Operational minimal | 0.9376 | 0.9553 | 0.0447 |
| Operational without derived outage severity | 0.9443 | 0.9636 | 0.0364 |
| Scenario descriptors only | 0.9596 | 0.8002 | 0.1998 |
| Hybrid reduced | 0.9560 | 0.8120 | 0.1880 |

For a planning application, these results motivate separate treatment of score ranking and the binary decision rule. A ranking can allocate a defined number of detailed assessments, whereas an alarm threshold determines which cases are flagged without fixing the resulting workload. The feature-set comparison therefore provides guidance on both information requirements and the intended use of the output. A transferred alarm threshold would be assessed against the target feeder's screening requirements. The demonstrated contribution is reuse of a screening model between IEEE-33 and IEEE-69 within the shared scenario-modelling framework. The following pruning analysis examines the practical consequence: how many low-voltage cases remain in the selected subset when the detailed-assessment budget is fixed.

### *4.4. Scenario Pruning and Computational Implications*

Figure 11 shows how undervoltage-case capture changes with the fraction of scenarios retained for detailed assessment. The principal comparison fixes this fraction at 30%, equivalent to 3000 cases in each 10,000-scenario library. Evaluating operational and descriptor-based rankings at the same budget makes their assessment requirements directly comparable. Within-network out-of-fold ranking and direct cross-network transfer then establish how effectively the selected subset concentrates undervoltage cases under each evaluation setting.

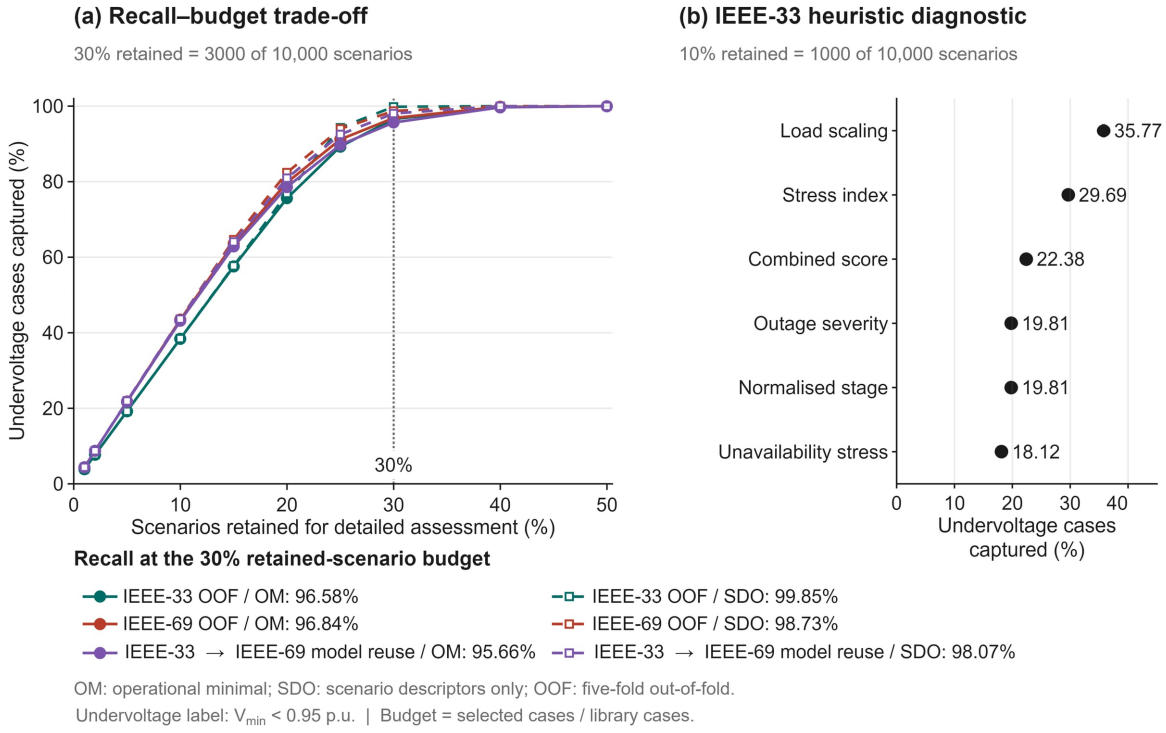



**Figure 11.** Recall–budget trade-off and heuristic ranking reference. (**a**) Undervoltage-case capture via the learned ridge-logistic rankings for both feeders and for model reuse across retained budgets; the principal budget is 30%. (**b**) Illustrative diagnostic of simple engineering rankings for IEEE-33 only at a 10% budget (1000 of 10,000 scenarios); no matched heuristic results are reported for IEEE-69 or at the 30% budget.

The pruning results show that low-voltage cases can be concentrated into a substantially smaller assessment subset. For IEEE-33, positive cases constitute 26% of the complete library. Retaining 30% therefore provides only four percentage points of additional capacity above the theoretical minimum fraction needed to capture every positive case. The achieved

capture at this budget indicates that the ranking places most low-voltage scenarios near the beginning of the assessment sequence.

The precision results provide a complementary interpretation. With 2600 positives and 3000 retained cases, the maximum possible precision is 86.67%. The descriptor-based precision of 86.53% approaches this ceiling, while the operational representation achieves 83.70%. These values explain the usefulness of the selected subset: a high proportion of the detailed evaluations is directed towards operating states that violate the chosen voltage criterion.

For renewable-integrated urban planning, the retained fraction makes the assessment workload explicit, while the missed-case count exposes the consequence of that allocation. The 30% result is therefore an evaluated operating point on a budget–capture relationship. Table 5 consolidates the within-network and transfer results at this common retained fraction.

**Table 5.** Undervoltage-case capture at a retained budget of 3000 out of 10,000 scenarios.

| Evaluation | Features | Positives | Captured | Recall (%) | Missed |
|---|---|---|---|---|---|
| IEEE-33 OOF | Operational | 2600 | 2511 | 96.58 | 89 |
| IEEE-33 OOF | Descriptors | 2600 | 2596 | 99.85 | 4 |
| IEEE-69 OOF | Operational | 2282 | 2210 | 96.84 | 72 |
| IEEE-69 OOF | Descriptors | 2282 | 2253 | 98.73 | 29 |
| IEEE-33 → IEEE-69 | Operational | 2282 | 2183 | 95.66 | 99 |
| IEEE-33 → IEEE-69 | Descriptors | 2282 | 2238 | 98.07 | 44 |

Operational = operational minimal (OM); descriptors = scenario descriptors only (SDO). All configurations use ridge-logistic regression. OOF denotes pooled five-fold out-of-fold scores. Undervoltage cases satisfy $V_{\min,s} < 0.95$ p.u. Recall@3000 is the captured count divided by the total undervoltage count; missed cases are undervoltage scenarios outside the retained subset. All rows retain 30% of the candidate scenarios, corresponding to 70% fewer cases selected for subsequent assessment.

The IEEE-69 results extend the retained-budget assessment to the larger radial benchmark. OM captures 2210 of the 2282 undervoltage cases, leaving 72 outside the selected subset. SDO captures 43 additional undervoltage cases at the same budget, reducing the missed count to 29. This improvement quantifies the value of scenario-construction information when the planning process already specifies the load multipliers and requested PV injection.

Cross-network application preserves most of the capture obtained by the corresponding IEEE-69 within-network rankings. Operational feature transfer captures 27 fewer undervoltage cases, a recall difference of 1.18 percentage points, while descriptor-based transfer captures 15 fewer, a difference of 0.66 percentage points. These descriptive comparisons demonstrate useful reuse of the screening relationship across the two benchmark feeders within the shared scenario-modelling framework. The practical contribution is an explicit choice between operational-input and scenario-informed prioritisation, with both evaluated against the same detailed-assessment budget.

Figure 11b is retained as a compact diagnostic of simple engineering rankings on IEEE-33 at a stringent 10% budget, corresponding to 1000 of 10,000 scenarios. At this budget, the six heuristic rankings captured between 18.12% and 35.77% of the IEEE-33 undervoltage cases, compared with the random-selection expectation of 10%. The principal learned-model evaluation remains the common 30% budget of Figure 11a and Table 5. Because no matched heuristic results are reported for IEEE-69 or at 30%, Figure 11b supports neither a learned-model advantage over heuristics on IEEE-69 nor a matched heuristic comparison at the principal budget.

Uniform random selection of 30% of the candidate scenarios would capture 30% of positive cases in expectation. The IEEE-33 operational ranking provides approximately 3.22 times this expected capture at the same assessment budget. This calculated enrichment describes the concentration achieved relative to an uninformative selection rule. Solver-call reduction and elapsed-time reduction are distinct. We let $t_{\mathrm{PF}}$ be the mean solver time per an AC power-flow evaluation, while $T_{\mathrm{score}}$ is the batch score-computation time, and $T_{\mathrm{other}}$ denotes the time for feature preparation, data checks, ranking, and any additional audit or communication overhead. Training and the initial labelled simulations are additional offline costs. For $N$ candidate scenarios of which $k$ are retained, the screening-stage time and the time required to solve every candidate are approximated by

$$T_{\mathrm{screen}} = T_{\mathrm{score}} + kt_{\mathrm{PF}} + T_{\mathrm{other}},\ T_{\mathrm{full}} \approx Nt_{\mathrm{PF}}. \tag{9}$$

In Equation (9), all $T$ terms are batch times (s), $t_{\mathrm{PF}}$ is the mean solver time per case (s), and $N$ and $k$ are case counts. The approximation assumes that the retained cases have representative mean solver cost; the individual retained-case runtimes and pipeline overheads were not recorded as a separate measurement.

Table 6 reports the timing evidence. The recorded mean AC power-flow solver times per scenario were 0.02222874 s for IEEE-33 and 0.02469705 s for IEEE-69. Multiplying these means by 10,000 gives solver-only estimates of 222.29 s and 246.97 s for the full libraries. At 3000 retained cases, the corresponding estimates are 66.69 s and 74.09 s, giving projected gross differences of 155.60 s and 172.88 s before screening overhead. Only the per-scenario means are recorded timing results; the batch totals and differences are projections that assume representative mean solver cost for the retained cases. Retained-subset runtimes and feature preparation, scoring, ranking, and other pipeline overheads were not measured separately. Exhaustive AC assessment therefore remains practical for these small benchmark feeders, and the timing evidence does not, by itself, establish a net computational advantage from machine learning. The 70% reduction concerns the number of cases selected for a subsequent assessment stage, not a measured end-to-end time saving. All development scenarios were already solved to obtain labels, so no saving is claimed for constructing the labelled libraries. Potential benefits for repeated studies, engineer review or more expensive downstream analyses require separate cost and timing measurements.

**Table 6.** Recorded per-scenario solver times and calculated solver-only batch estimates.

| Quantity | IEEE-33 | IEEE-69 |
|---|---|---|
| Recorded mean solver time per scenario (s) | 0.02222874 | 0.02469705 |
| Full-library solver time, 10,000 cases (s; calculated) | 222.29 | 246.97 |
| Retained-subset solver time, 3000 cases (s; calculated) | 66.69 | 74.09 |
| Gross solver-only difference, 7000 cases (s; calculated) | 155.60 | 172.88 |

Only the per-scenario means are recorded timing results. Batch totals and differences are products of those means and the stated case counts, calculated from the unrounded means and rounded to two decimals; they exclude scoring, preparation, ranking, and other overheads. The initial labelled libraries were solved in full, so these estimates concern a subsequent screening stage, not savings during dataset construction.

### *4.5. Interpretation of Feature Importance*

Figure 12 shows the final operational permutation-importance plots. Residual demand and PV shape dominate the displayed PR-AUC decrease in both networks. Generation unavailability contributes less, while national stage and derived outage severity show comparatively small individual decreases. The result supports a demand-and-PV-dominated ranking mechanism within the constructed scenario family.

The prominence of residual demand and PV shape is consistent with their roles in the scenario equations. Residual demand influences the loading multiplier, while the PV shape contributes to both the load mapping and the prescribed active-power injection. Their importance therefore connects the learned ranking to identifiable demand–generation relationships within the constructed operating states. This interpretation provides a practical basis for examining screening inputs: demand estimates and PV availability deserve particular attention when preparing candidate scenarios. Generation-unavailability and load-shedding indicators supply additional operating context, with smaller individual permutation effects in the displayed analysis. Because some predictors are strongly dependent, the importance values describe model reliance under the adopted representation; they do not partition independent causal effects. The resulting explanation nevertheless allows the principal screening inputs to be traced back to the physical scenario construction.

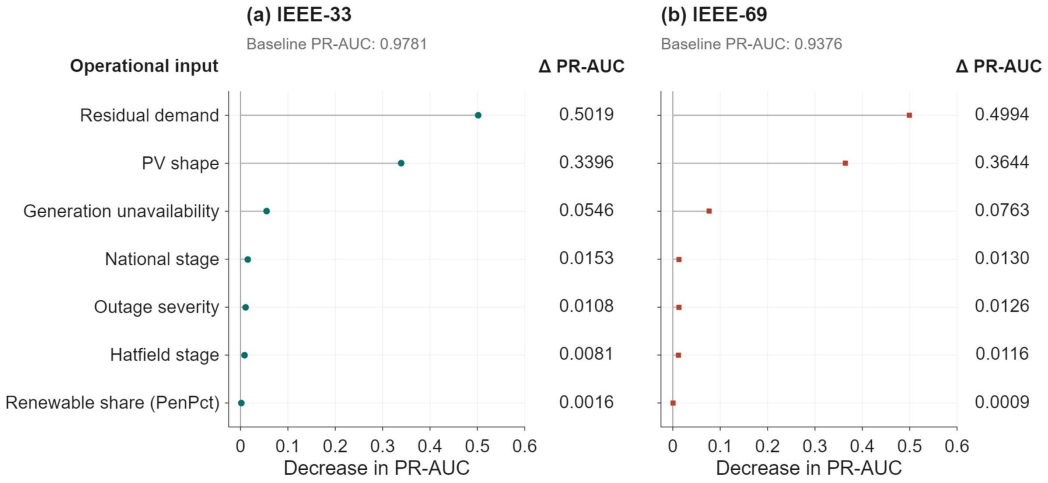



**Figure 12.** Operational feature permutation importance.

### *4.6. Prospective Application to Urban Network Planning*

The demonstrated contribution is a benchmark assessment process that connects renewable-demand operating conditions to the allocation of detailed distribution-network studies. Historical operational context informs the candidate scenarios, AC power flow supplies the electrical reference, and the learned ranking identifies a prioritised subset for further examination. Urban and smart-city network planning is a promising, but not yet demonstrated, area of application. In that setting, the workflow would address a practical planning question: which combinations of demand, photovoltaic availability, and supply-scarcity conditions warrant a detailed voltage assessment within the available study budget?

For municipal utilities, campus operators and industrial energy managers, the choice of screening configuration would depend on whether contextual operating indicators alone or the full scenario-construction inputs are available during study preparation. A future application would need to retain scenario identifiers, input assumptions, model version, screening scores, and the selected assessment budget and validate the input mapping, threshold, and budget on a feeder-specific model with measured data. Detailed solutions for the prioritised cases could then inform subsequent investigations of voltage support, storage scheduling, or flexible demand. No municipal deployment, measured urban-feeder validation, or smart-city system integration is claimed in this study.

### *4.7. Applicability and Further Validation*

The evaluation establishes retrospective undervoltage-screening performance for two balanced radial benchmarks under historically informed, stress-enriched scenario envelopes.

The reported positive fractions are properties of the constructed libraries, not measured probabilities of feeder events. Both datasets use a shared historical source and related scenario mappings; random scenario splits therefore do not establish independence of operating periods. The results are conditional on the prescribed photovoltaic active-power injections, unchanged inverter reactive-power contribution, fixed topology and adopted voltage threshold. Upper-voltage compliance, performance with voltage-regulating inverters, sub-hourly solar ramps and dynamic recovery have not been evaluated. The missed-case counts in Table 5 also show that the retained subset is not exhaustive. A feeder-specific application requires independent data and an assessment budget chosen against both the frequency and severity of missed cases.

Reproducibility should be supported by a versioned archive containing the final data definitions, network cases, scenario settings, evaluation partitions, and numerical outputs. These provisions establish a clear route from the present benchmark assessment to feeder-specific planning use.

## 5. Conclusions

This study evaluated a digital engineering framework for steady-state undervoltage screening and scenario prioritisation in renewable-integrated distribution planning, to support city network planners in the transition towards smarter, cleaner, and more sustainable cities. By way of illustrative example, South African operational indicators, Tshwane/Hatfield load-shedding histories, and modelled photovoltaic profiles were linked to prescribed network operating states, with the AC power flow providing a consistent undervoltage label at $V_{\min} < 0.95$ p.u. The evaluation comprised 10,000 converged scenarios for each of the IEEE-33 and IEEE-69 feeders. The operational feature model achieved mean PR-AUC values of 0.9782 and 0.9476 within IEEE-33 and IEEE-69, respectively. Applying the IEEE-33-trained model to IEEE-69 achieved PR-AUC 0.9376 and classification recall 0.9553. This result demonstrates model reuse within a shared scenario-modelling framework: both libraries use the same historical operational source and related generation rules, and the IEEE-69 envelope was adjusted during pilot design. It does not establish generalisation to an independently sampled operating distribution, a future period or a real feeder.

At a retained budget of 3000 out of 10,000 scenarios, operational feature ranking captured 96.58% of IEEE-33 undervoltage cases, 96.84% of IEEE-69 cases and 95.66% under model reuse. The scenario-descriptors-only configuration captured 2596 of 2600 IEEE-33 undervoltage cases (99.85%). That result depends on known scenario-construction inputs and is conditional on the scenario generator, stress-enriched sampling, and prescribed benchmark networks; the SDO results cannot be read as screening indicators on real feeders. Neither feature representation has been validated with field data.

Retaining 30% of the candidate library selects 70% fewer cases for subsequent assessment. The recorded mean solver times (Table 6) imply projected solver-only differences of 155.60 s and 172.88 s per 10,000-scenario batch before screening overhead. These are projections rather than measured end-to-end savings, and they do not establish compliance of the deferred scenarios.

The combined results establish how predictor availability, cross-network model reuse and assessment budget influence the concentration of low-voltage cases in a prioritised scenario subset. This provides an explicit connection between machine-learning performance and an engineering planning decision. The principal demand and PV predictors also give the ranking an interpretable relationship with the scenario construction.

Urban and smart-city network planning is a prospective application of this work. The framework offers a benchmark workflow for prioritising operating scenarios for feeder-specific analysis. Its present evidence concerns balanced, fixed-topology networks with

prescribed active and reactive injections. It does not establish dynamic stability, complete voltage compliance, optimal investment decisions, or operational readiness. Independent feeder measurements and implementation-specific assessments will be required its before application to an urban distribution system.

Future work should prioritise validation with measured feeder demand, photovoltaic output, and voltage records using chronological or source-record-grouped testing. Further extensions should evaluate upper-voltage events, inverter reactive-power controls, and the severity of undervoltage cases missed at each retained-case budget. Implementation studies should measure end-to-end costs and latency, including feature preparation, model scoring, and downstream engineering assessments. Dynamic stability would require a separate set of disturbance models, system dynamics, and stability criteria.

**Author Contributions:** Conceptualisation, R.W., M.S., G.A., M.A.-G. and R.N.; methodology, R.W., M.S., G.A. and M.A.-G.; software, R.W., M.S. and G.A.; validation, R.W., M.S., G.A., M.A.-G. and R.N.; formal analysis, R.W., M.S. and G.A.; investigation, R.W., M.S. and G.A.; resources, M.S., G.A. and M.A.-G.; data curation, R.W., M.S. and G.A.; writing—original draft preparation, R.W., M.S. and G.A.; writing—review and editing, R.W., M.S., G.A. and M.A.-G.; visualisation, R.W., M.S., G.A. and M.A.-G.; supervision, M.S., G.A., M.A.-G. and R.N.; project administration, M.S., G.A., M.A.-G. and R.N.; All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by Teesside University through a PhD scholarship awarded to R.W.

**Data Availability Statement:** The public Tshwane load-shedding dataset is available from the University of Pretoria Research Data Repository [8]. The other source data and modelled photovoltaic inputs are identified in Table 1. The derived benchmark scenario data, analysis code, and figure-generation records supporting the reported results are available from the corresponding author upon reasonable request, subject to the original source licences.

**Acknowledgments:** The authors are grateful to Teesside University, UK, for funding the PhD scholarship of Rajitha Wattegama. During the preparation of this manuscript, the authors used ChatGPT 6.1 Sol (OpenAI) and Claude Opus 5.5 (Anthropic) to assist when checking the results and during revision to support drafting and preparation of mathematical notations. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Notation and Abbreviations

| Symbol | Definition and Units |
|---|---|
| $V_{i,s}$ | Voltage magnitude at bus $i$ in scenario $s$ (p.u.). |
| $V_{\min,s}$ | Minimum bus-voltage magnitude in scenario $s$ (p.u.). |
| $\tau$ | Voltage-screening threshold (p.u.). |
| $P_i^0$ | Base active-power demand at bus $i$ (MW). |
| $Q_i^0$ | Base reactive-power demand at bus $i$ (MVAr). |
| $A_{i,s}$ | Allocated PV active power at bus $i$ in scenario $s$ (MW). |
| $p_i$, $q_i$ | Net active- and reactive-power injections, respectively, expressed on the case MVA base (p.u.). |
| $\delta_i$ | Voltage angle at bus $i$ (rad). |
| $G_{ij}$, $B_{ij}$ | Admittance-matrix entries (p.u.). The matrix entry $G_{ij}$ is distinct from the source indicator $G$. |
| $G$ | Source generation-unavailability indicator. |
| $N$ | Number of eligible scenarios. |
| $N_+$ | Number of positive-label scenarios. |
| $k$ | Number of scenarios retained for detailed assessment. |
| TP | True-positive count. |
| FN | False-negative count. |
| FP | False-positive count. |
| TN | True-negative count. |

| | |
|---|---|
| $L$, $Q$ | Active-demand multiplier and additional reactive-demand modifier, respectively (dimensionless). |
| $f$ | Local PV fraction (dimensionless). |
| $S$ | Composite stress descriptor (dimensionless). |
| $\eta$ | Load-heterogeneity multiplier (dimensionless). |
| $r$, $g$, $e$, $\ell$ | Normalised residual demand, generation unavailability, renewable share, and manual load reduction, respectively (dimensionless). |
| $n$, $m$, $h$ | Normalised national, municipal, and Hatfield stages, respectively (stage number divided by eight; dimensionless). |
| $v$ | Normalised PV output (dimensionless). |
| $a$ | Active-outage flag. |
| $o$ | Derived outage-severity indicator (dimensionless). |
| $t_{\mathrm{PF}}$ | Mean AC power-flow solver time per case (s). |
| $\mathcal{B}$; $i$, $j$; $s$, $t$ | Set of bus identifiers; bus indices; scenario index; and historical-record index, respectively. |
| $y_s$; $\mathbf{1}\{\cdot\}$ | Binary undervoltage label; indicator operator (dimensionless). |
| $\mathrm{TP}_k$ | Positive scenarios within the retained subset (count). |
| $R_{\mathrm{calls}}$ | Fractional reduction in cases selected for solver-based follow-up (dimensionless). |
| $q_{0.05,x}$; $q_{0.95,x}$; $z_x$ | Fifth and ninety-fifth percentiles, in the units of source variable $x$; normalised value (dimensionless). |
| $w_t$; $H_t$ | Sampling weight (dimensionless); unnormalised Hatfield stage (stage number). |
| $\rho_{\mathrm{RE}}$ | Ratio of national renewable generation to residual demand (%). |
| $P^{\mathrm{net}}_{i,s}$; $Q^{\mathrm{d}}_{i,s}$ | Net active demand before export clipping (MW); prescribed reactive demand (MVAr). |
| $P_{\mathrm{PV,requested}}$ | Requested aggregate photovoltaic active power (MW). |
| $\varepsilon_L$, $\varepsilon_Q$, $\varepsilon_f$ | Zero-mean Gaussian scenario noise terms with standard deviations $\sigma_L$, $\sigma_Q$ and $\sigma_f$ (dimensionless). |
| $\xi_i$; $d_i$ | Standard-normal draw; bus index divided by the maximum bus index (dimensionless). |
| $T_{\mathrm{score}}$, $T_{\mathrm{other}}$, $T_{\mathrm{screen}}$, $T_{\mathrm{full}}$ | Batch-time terms in Equation (9) (s). |
| PQ bus; voltage-controlled PV bus | Power-flow bus types with prescribed active and reactive power, or with prescribed active power and voltage magnitude. |
| OM, ON, SDO, HR | Feature sets of operational minimal; operational without derived outage severity; scenario descriptors only; hybrid reduced. |
| AC | Alternating current. |
| AUC | Area under a curve. |
| ML | Machine learning. |
| OOF | Out of fold. |
| PR | Precision–recall. |
| PV | Photovoltaic generation—the power-flow bus type is always written "voltage-controlled PV bus". |
| ROC | Receiver operating characteristic. |

## Appendix A. Scenario-Generation Equations

The scenario mappings use the dimensionless indicators defined in Section 3.2. Here, $e$ is the normalised renewable share ($\rho_{\mathrm{RE}}$), $\ell$ is normalised manual load reduction, and $v$ is normalised photovoltaic output. Subscripts 33 and 69 identify the feeder-specific mappings. The coefficients and clipping bounds are prescribed scenario-design choices. The noise terms $\varepsilon_L$, $\varepsilon_Q$ and $\varepsilon_f$ are zero-mean Gaussian with the standard deviations $\sigma_L$, $\sigma_Q$ and $\sigma_f$ in Table 2. For the IEEE-33 feeder, the active-load multiplier is given by

$$L_{33} = \mathrm{clip}(0.34 + 0.34r + 0.08g + 0.06n + 0.05m + 0.08h + 0.06o + 0.02\ell - 0.08v + \varepsilon_L,\ 0.28,\ 1.08). \tag{A1}$$

For the IEEE-69 feeder, the corresponding multiplier is

$$L_{69} = \mathrm{clip}(0.32 + 0.32r + 0.075g + 0.055n + 0.045m + 0.070h + 0.050o + 0.018\ell - 0.085v + \varepsilon_L,\ 0.26,\ 1.03). \tag{A2}$$

Positive stress terms in Equations (A1) and (A2) increase the assumed load envelope, while the negative photovoltaic-shape term reduces it. The bounds are scenario assumptions rather than fitted feeder operating limits. For the IEEE-33 feeder, the additional reactive-demand modifier is given by

$$Q_{33} = \mathrm{clip}\big(0.82 + 0.040h + 0.035g + 0.025o + \varepsilon_Q,\ 0.76,\ 1.04\big). \tag{A3}$$

For the IEEE-69 feeder, the corresponding modifier is

$$Q_{69} = \mathrm{clip}\big(0.80 + 0.035h + 0.030g + 0.020o + \varepsilon_Q,\ 0.74,\ 1.00\big). \tag{A4}$$

The total reactive-demand scale in Equation (7) is the product of the modifier in Equations (A3) and (A4) and the active-load multiplier. The raw local photovoltaic fractions of the IEEE-33 and IEEE-69 feeders combine the normalised renewable share, the sampled photovoltaic output, and a Gaussian perturbation, which are given by

$$\begin{aligned} f_{33,\mathrm{raw}} &= 0.06 + 0.58e + 0.22v + \varepsilon_f, \\ f_{69,\mathrm{raw}} &= 0.065 + 0.60e + 0.23v + \varepsilon_f. \end{aligned} \tag{A5}$$

The local fraction used in the network model is obtained by the clipping operation stated after Equation (A6). The requested aggregate photovoltaic active power is the product of the total base active demand, the local photovoltaic fraction, and the normalised photovoltaic output:

$$P_{\mathrm{PV,requested}} = \Big(\sum\nolimits_i P_i^0\Big) f\, v. \tag{A6}$$

In Equation (A6), $f = \mathrm{clip}(f_{\mathrm{raw}})$ is the clipped form of the corresponding raw expression in Equation (A5); the bounds are those of the archived scenario-generation code.

The requested power is an instantaneous active-power quantity in MW, not energy. Effective injection can be lower after the per-connection export cap. The composite stress descriptor is a clipped, dimensionless weighted combination of the historical indicators, given by

$$S = \mathrm{clip}(0.30r + 0.18g + 0.13n + 0.10m + 0.15h + 0.09o + 0.04\ell - 0.04v,\ 0,\ 1). \tag{A7}$$

$S$ is an input descriptor and not a stability index or a risk label.

The heterogeneity factor in Equation (7) is $\eta_{i,s} = \mathrm{clip}\ (1 + 0.08\xi_i + 0.10hd_i + 0.06gd_i)$, where $\xi_i$ is a standard-normal draw, $d_i$ is the bus index divided by the maximum bus index, and $h$ and $g$ are the normalised Hatfield stage and generation unavailability of the scenario.

## References


1. Máslo, K.; Kasembe, A. Mitigation measures for photovoltaics retrofit. *IEEE Trans. Sustain. Energy* **2017**, *9*, 333–339. [CrossRef]
2. Kabeyi, M.J.B.; Olanrewaju, O.A. Sustainable energy transition for renewable and low carbon grid electricity generation and supply. *Front. Energy Res.* **2022**, *9*, 743114. [CrossRef]
3. Farrokhabadi, M.; Cañizares, C.A.; Simpson-Porco, J.W.; Nasr, E.; Fan, L.; Mendoza-Araya, P.A.; Tonkoski, R.; Tamrakar, U.; Hatziargyriou, N.D.; Lagos, D.; et al. Microgrid stability definitions, analysis, and examples. *IEEE Trans. Power Syst.* **2019**, *35*, 13–29. [CrossRef]
4. Ning, J.; Li, F.; Wan, Y.; Liu, Y. Enhancing stability in multimode hybrid power supply systems under stepped constant power loads: A synergetic control approach. *Int. J. Circuit Theory Appl.* **2026**. *Epub ahead of print*. [CrossRef]
5. Adebiyi, A.A.; Moloi, K. Renewable energy source utilization progress in South Africa: A review. *Energies* **2024**, *17*, 3487. [CrossRef]
6. Majenge, L.; Mpungose, S.; Msomi, S. Threshold effects on South Africa's renewable energy–economic growth–carbon dioxide emissions nexus: A nonlinear analysis using threshold-switching dynamic models. *Energies* **2025**, *18*, 4642. [CrossRef]
7. Wattegama, R.; Short, M.; Aggarwal, G.; Al-Greer, M.; Naidoo, R. A systematic review of hierarchical control frameworks in resilient microgrids: South Africa focus. *Energies* **2026**, *19*, 644. [CrossRef]
8. Twala, S. *Historical Load Shedding Data for the City of Tshwane, 2015 to 2023 [Dataset]*; University of Pretoria Research Data Repository: Pretoria, South Africa, 2024. Available online: https://researchdata.up.ac.za/articles/dataset/Historical_load_shedding_data_for_the_City_of_Tshwane_2015_to_2023/27980558 (accessed on 14 September 2026).
9. Eskom. *Weekly System Status Reports [Internet]*; Eskom: Johannesburg, South Africa, 2026. Available online: https://www.eskom.co.za/eskom-divisions/tx/system-adequacy-reports/ (accessed on 14 September 2026).
10. Wu, C. A comparative study of PVWatts and EnergyPlus in simulating electricity production of fixed PV arrays. *Clean. Energy Syst.* **2026**, *14*, 100259. [CrossRef]

11. Huang, B.; Wang, J. Applications of physics-informed neural networks in power systems: A review. *IEEE Trans. Power Syst.* **2022**, *38*, 572–588. [CrossRef]
12. Babaali, A.H.; Ameli, M.T. Multi-class imbalanced learning for short-term voltage stability assessment. *E-Prime—Adv. Electr. Eng. Electron. Energy* **2025**, *14*, 101128. [CrossRef]
13. Mossie, M.A.; Yetayew, T.T.; Bitew, G.T.; Yenealem, M.G.; Beza, T.M. Machine learning algorithms for voltage stability assessment in electrical distribution systems. *Sci. Rep.* **2025**, *15*, 31980. [CrossRef] [PubMed]
14. Wiese, M.; van der Westhuizen, L.-M. Impact of planned power outages (load shedding) on consumers in developing countries: Evidence from South Africa. *Energy Policy* **2024**, *187*, 114033. [CrossRef]
15. Zimmerman, R.D.; Murillo-Sánchez, C.E. *MATPOWER User's Manual [Internet]*, version 8.0; Zenodo: Geneva, Switzerland, 2024. [CrossRef]
16. Baran, M.E.; Wu, F.F. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Trans. Power Deliv.* **1989**, *4*, 1401–1407. [CrossRef]
17. Baran, M.E.; Wu, F.F. Optimal capacitor placement on radial distribution systems. *IEEE Trans. Power Deliv.* **1989**, *4*, 725–734. [CrossRef]
18. Wu, Y.; Sicard, B.; Gadsden, S.A. A review of physics-informed machine learning methods with applications to condition monitoring and anomaly detection. *arXiv* **2024**, arXiv:2401.11860. [CrossRef]
19. National Laboratory of the Rockies. *PVWatts Calculator [Internet]*; National Laboratory of the Rockies: Golden, CO, USA, 2026. Available online: https://pvwatts.nlr.gov/ (accessed on 14 September 2026).
20. Dobos, A.P. *PVWatts Version 5 Manual*; 2014 Sep. Report No.: NREL/TP-6A20-62641; National Renewable Energy Laboratory: Golden, CO, USA, 2014. Available online: https://www.osti.gov/biblio/1158421 (accessed on 14 September 2026).
21. Trinh, P.-H.; Chung, I.-Y. Integrated active and reactive power control methods for distributed energy resources in distribution systems for enhancing hosting capacity. *Energies* **2024**, *17*, 1642. [CrossRef]
22. Zimmerman, R.D.; Murillo-Sánchez, C.E.; Thomas, R.J. MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Trans. Power Syst.* **2011**, *26*, 12–19. [CrossRef]
23. MathWorks. *Fitclinear: Fit Binary Linear Classifier to High-Dimensional Data [Internet]*; Statistics and Machine Learning Toolbox Documentation; The MathWorks, Inc.: Natick, MA, USA, 2026. Available online: https://www.mathworks.com/help/stats/fitclinear.html (accessed on 28 October 2026).
24. Cerulli, G. Model selection and regularization. In *Fundamentals of Supervised Machine Learning: With Applications in Python, R, and Stata*; Statistics and Computing; Springer: Cham, Switzerland, 2023; pp. 59–146. [CrossRef]
25. Davis, J.; Goadrich, M. The relationship between precision-recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, USA, 25–29 June 2006*; Association for Computing Machinery: New York, NY, USA, 2006; pp. 233–240. [CrossRef]
26. Strobl, C.; Boulesteix, A.-L.; Kneib, T.; Augustin, T.; Zeileis, A. Conditional variable importance for random forests. *BMC Bioinform.* **2008**, *9*, 307. [CrossRef] [PubMed]


**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.