UNIVERSITY OF WEST ATTICA

eISSN:3136-7658

**Applied Manufacturing Intelligence Journal**

Vol. 5, No. 3, July, 2026




# Constraint Screened Microgrid Black Start Strategy Generation Using ConvNeXtV2 CBAM

Marcel Jasiński[1,*], Władysław Gustaw Jach[1], Marcel Knapik[1], Dominik Błasik[1]

[1] *Faculty of Electrical and Control Engineering, Silesian University of Technology, Gliwice, 44-100, Poland*

Received 24 March 2026; Revised 30 June 2026; Accepted 03 July 2026; Published 09 July 2026

**Abstract**

Microgrid black start requires coordinated inverter energization and staged load pickup under limited current capability, uncertain demand recovery, and incomplete measurements. This paper develops a ConvNeXtV2 based strategy generator with a convolutional block attention module to encode inverter histories and load restoration states, propose switching actions, and submit each proposal to an independent feasibility screen. Representation learning, sequence generation, and operational acceptance are separated so that physical constraints are evaluated independently of proposal probability. The evaluation covers 5100 operating records, including 3600 training, 900 validation, and 600 testing records, and compares six supervisory methods on common testing identifiers. The full configuration achieves 576 accepted episodes out of 600, compared with 558 out of 600 for the ConvNeXtV2 encoder baseline. Mean restoration durations within the respective accepted subsets are 98 and 106 seconds. The corresponding improvement is three percentage points in acceptance and eight seconds in restoration duration. Fifteen diagnostic panels examine operating coverage, transient response, duration distributions, computational tradeoffs, observation loss, and pickup sequencing. Joint analysis of voltage and current margins identifies cases in which satisfactory voltage recovery coexists with excessive transient current. The proposed framework connects heterogeneous restoration records to executable supervisory decisions through an explicit candidate acceptance interface. Its operating design supports the assessment of service recovery, constraint satisfaction, and decision latency within a common restoration procedure.

*Keywords:* Microgrid Black Start, Inverter Restoration Data, Load Pickup Sequencing, ConvNeXtV2 CBAM, Constraint Screened Generation

## 1. Introduction

Inverter based black start depends on coordinating reference establishment, resource energization, and the progressive reconnection of interrupted demand [1]. Grid forming restoration studies examine how locally established voltage and frequency references support the formation and expansion of an energized island [2]. Current limiting research identifies converter saturation as a source of changes in the electrical response during energization [3]. Load pickup studies further relate interruption history and load composition to the transient demand imposed on the available sources [4]. A restoration controller must therefore distinguish continuous power adequacy from the short duration capability needed for transformer energization and motor acceleration. Selecting the pickup time requires both the eventual steady consumption and the expected transient response of the disconnected block. The supervisory task comprises device and load selection, dwell interval determination, and observation-based authorization of subsequent actions. These decisions unfold while storage energy declines and while the electrical characteristics of the energized island change. A sequence that appears satisfactory under one ordering can become unacceptable after a different load is connected first. The

* Corresponding author: marcel.jas@polsl.pl

https://doi.org/10.68406/amij.2026.vol5iss3np1:1-19

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

resulting problem combines discrete topology changes, continuous operating limits, and uncertain transient demand. It also requires a clear boundary between decisions made by a supervisory algorithm and the fast regulation performed inside individual converters. Explicit coordination between supervisory decisions and local regulation is required to keep switching schedules within the current support capability of installed devices. The engineering objective is to restore priority demand through actions that remain interpretable and admissible as the energized island develops.

Restoration records contain information at several incompatible scales. Inverter measurements describe voltage formation, active and reactive support, current clipping, and the availability of energy reserves. Load records describe interrupted demand, priority, previous pickup attempts, and the time since disconnection. Work on restoration observability emphasizes the importance of associating such records with the topology and status of the recovering network [5]. Data driven restoration studies have explored representations that reduce heterogeneous measurements to compact decision states [6]. Preserving temporal signatures helps distinguish a tolerable pickup from the onset of a protection operation. Research on cold load pickup further indicates that demand recovery depends on interruption duration and load composition [7]. The same nominal load block may therefore require different waiting times under different initial conditions. Delayed or missing inverter telemetry requires explicit quality indicators to distinguish unavailable observations from measured zero current. Methods for control with imperfect observations motivate explicit treatment of measurement quality and timestamp alignment [8]. A useful learning representation must retain these distinctions and construct each decision input exclusively from information available before the action. The representation must identify channels describing resource capability, demand, and data quality. The representation must preserve the relationship between the information available before an action, the selected operation, and the electrical response observed afterward. Preserving this relationship supports an unbiased assessment of policy generalization to unseen restoration episodes.

Xu et al. [9] formulate adaptive robust load restoration through coordinated distribution network reconfiguration and mobile energy storage. Their explicit constraints provide a useful engineering reference, although detailed models can increase the computational cost of repeated decisions. Learning based policies offer an alternative route for mapping observations to actions when a representative training collection is available [10]. Their limitations become more consequential during black start because an erroneous action can remove the voltage reference required to sustain the remaining energized network. Woo et al. [11] introduce ConvNeXt V2, a hierarchical convolutional architecture that incorporates global response normalization. Convolutional block attention separately refines channel and spatial responses, suggesting a possible mechanism for emphasizing transient intervals and measurement groups [12]. Applying these components to electrical restoration requires adapting their input geometry and validation protocol to equipment histories. The device axis requires a documented physical ordering so that adjacent feature positions have a consistent engineering interpretation. Research on constrained learning supports separating a learned proposal from an admissibility mechanism [13]. Calibration research in medical imaging relates prediction confidence to observed outcome frequencies; its relevance here is methodological, with operational calibration requiring validation on restoration outcomes [14]. The research gap is therefore a coherent interface joining restoration records, a modern convolutional encoder, a sequence generator, and an external acceptance process. Such an interface should record accepted actions, rejected proposals, and recovery behavior with comparable detail because each contributes to the operational burden of the controller.

Motivated by the need to connect heterogeneous operating histories with executable restoration decisions, this study develops a ConvNeXtV2 CBAM framework that combines temporal feature encoding, staged action generation, and independent assessment of inverter and load pickup constraints. The model receives aligned inverter histories and load restoration descriptors, generates the next supervisory action, and passes the candidate through a separate screen before any execution is permitted. The contribution concerns the integration of aligned restoration histories, a ConvNeXtV2 CBAM representation, structured action generation, and an independent acceptance interface. The architectural components retain their established operations; the task-specific formulation defines their observation geometry, command semantics, and connection to electrical screening. The evaluation examines the resulting supervisory configuration through separate measures of service recovery, transient margins, and decision effort. The evaluation retains separate metrics for proposal admissibility, completed service, and the computational effort required to reach each restoration decision. The remaining sections follow this progression. Section 2 examines restoration scheduling, data representations, and learning based decision mechanisms. Section 3 defines the supervisory problem, the observation tensor, and the interfaces required for admissible actions. Section 4 describes the encoder, attention adaptation, decoder, training objectives, and operational screening procedure. Section 5 presents the evaluation protocol, two supporting tables, and analyses of operating coverage, transient response, comparative performance, sensitivity, and sequencing. Section 6 concludes the paper by addressing the methodological contribution, its present limitations, and the validation work required before engineering deployment.

## 2. Related Work

### 2.1. Restoration Scheduling and Inverter Capability

Restoration scheduling has commonly been expressed as a constrained selection of energization and load pickup actions. Research on distribution restoration uses switching decisions to coordinate network connectivity with the availability of distributed resources [15]. Such formulations make action dependencies explicit and can represent priority differences between critical and deferrable demand. Their engineering value depends on the fidelity of the electrical constraints and the correspondence between the optimization time step and the relevant transient. Steady operating point constraints address sustained overloads, while transient constraints address brief current surges. Pickup rules that account for available transient support can identify opportunities for earlier energization. Inverter based black start adds the need to establish a reference and to maintain a feasible sharing arrangement as additional resources join the island. Work on grid forming restoration examines coordination among converters and switching devices under these conditions [16]. The associated supervisory problem combines load priority with resource capability and energization dependencies. A high priority block may need supporting resources, a voltage ramp, or a different energization path before pickup is permissible. Current limiting behavior introduces another modeling issue because controller saturation can alter the dynamic response assumed by an unconstrained source model [17]. Research on transformer and motor energization connects restoration sequencing with magnetic and electromechanical transients [18]. These studies motivate constraints that depend on the type and state of the equipment being energized. Nevertheless, optimization formulations differ in how they represent inrush, protection timing, and uncertainty in recovered demand. Comparable scheduling assessments require consistent dynamic assumptions and a common definition of failed attempts. A meaningful assessment should report what each method is allowed to observe and which acceptance rules are applied after a proposed action. Interpreting differences in restoration duration consequently requires a common feasibility model and consistent outcome definitions.

### 2.2. Learning Representations for Restoration Records

Learning approaches replace some explicit decision calculations with a mapping estimated from previously generated or recorded episodes. Studies of learning-based restoration policies investigate this mapping under a range of network configurations and demand conditions [19]. The representation of an episode is consequential because the restoration outcome depends on both the present state and the path used to reach it. Restoration histories provide information about inverter recovery from clipping and unsuccessful load acceleration that complements current power values. Temporal encoders retain part of this history, but their input construction still requires careful alignment of device signals and action timestamps. Research on multivariate time series classification discusses the benefits and limitations of convolutional representations for structured temporal observations [20]. In restoration, the second dimension of such a representation may contain devices, features, or physically grouped channels. A documented ordering gives convolutions a consistent relationship to physical equipment and measurement groups. Missing signals require equally explicit treatment. An imputed value and an observed value can have identical numerical magnitude while conveying different operational confidence. Studies of missing data in learned control motivate explicit quality masks and observation age descriptors alongside imputed values [21]. Another difficulty arises when neighboring time windows from the same episode are divided between training and testing. Shared trajectories introduce dependence between the training and test sets and can inflate estimated generalization performance. Restoration data should therefore be partitioned at the level of scenario families, trajectories, or network realizations appropriate to the claimed generalization. The literature provides several useful encoder families whose evaluation also requires explicit controls for these data issues. Evidence about robustness requires an evaluation collection in which topology changes, resource conditions, and missingness patterns are varied independently enough to expose the intended failure modes.

### 2.3. Attention Mechanisms and Constrained Decisions

Hierarchical convolutional models offer increasing receptive fields and compact latent representations, which can be useful when restoration records contain both local transient features and longer operating trends. ConvNeXt V2 couples its architecture with a masked representation learning framework and global response normalization [22]. Those components originate in visual recognition, so their adaptation to restoration should preserve the mathematical operation while revisiting the physical meaning of the input arrangement. Channel attention can emphasize measurement families that carry decision relevant evidence, whereas spatial attention can highlight regions within the encoded history. The sequential channel and spatial refinement used by convolutional block attention provides a concrete mechanism for this selection [23]. Attention weights describe internal feature responses and require complementary intervention tests to assess the contribution of correlated sensor channels. Output semantics determine the operational scope of the learned model. Classification selects a predefined label, while restoration generation specifies a device, an operation, and a dwell interval under a changing feasible action set. Research on safety filters and constrained policy execution motivates an independent mechanism that can reject a learned candidate before it is submitted for physical execution [24]. The filter

specification requires declared modeling assumptions, a valid observation state, and a defined fallback for an empty admissible candidate set. The relevant gap lies between representation quality and executable supervision. A restoration method needs an explicit account of candidate generation, constraint evaluation, rejection logging, and state updates after execution. The present formulation uses this gap to organize its design. Comparisons between encoder variants are formulated as hypotheses requiring common data, matched training effort, and independent dynamic validation.

## 3. Problem Formulation and Restoration Data Representation

### 3.1. Supervisory State and Action Semantics

The target system is an islandable microgrid with at least one storage supported grid forming inverter, additional controllable resources, sectionalizing switches, and independently restorable load blocks. The supervisory controller operates at a higher control level than the inverter regulation loops. Local inverter loops retain responsibility for modulation, regulation, and current limiting, while the installed protection system retains its assigned authority. The supervisory controller requests an allowed resource transition, selects a load block for pickup, opens an admissible switch when recovery requires isolation, or schedules additional settling. An action record contains the operation type, target identifier, optional support setpoint, and minimum observation interval. These fields must be validated against a device registry before electrical admissibility is considered. Command eligibility requires the switch identifier to belong to the active topology before predicted power is assessed. The state includes energized connectivity, resource availability, storage reserve, restored load status, and measurement quality. Initial black start is represented explicitly by an empty energized island and an available reference resource. Later steps operate on the connected island established by earlier accepted actions.

Candidate sequences are ranked by minimizing a finite horizon cost comprising weighted unserved energy, switching operations, and rejected proposals. Electrical admissibility is enforced separately from the cost used to rank candidate sequences. Each priority weight is fixed before evaluation and attached to a specified load block. Unserved energy is measured over the restoration horizon, whereas switching and rejection counts are dimensionless quantities with declared penalty weights. The following equation combines these contributions into the sequence ranking objective:

$$J = \sum_{k} \Delta t_k \sum_{j} w_j P_j \left(1 - z_{jk}\right) + \lambda_s N_s + \lambda_r N_r \tag{1}$$

Here, $z_{jk}$ denotes the served status of load block $j$ at decision $k$, $P_j$ its requested power, and $w_j$ its priority weight. The counts $N_s$ and $N_r$ refer to switching operations and rejected proposals, respectively. The time interval $\Delta t_k$ is measured in hours when the first term is expressed as weighted kilowatt hours. The penalty coefficients then have the corresponding cost units per event. Priority weights are dimensionless and nonnegative, and the event penalties are nonnegative costs per switching operation or rejected proposal. Minimizing the objective favors earlier delivery of higher-priority demand while accounting for the declared event penalties.

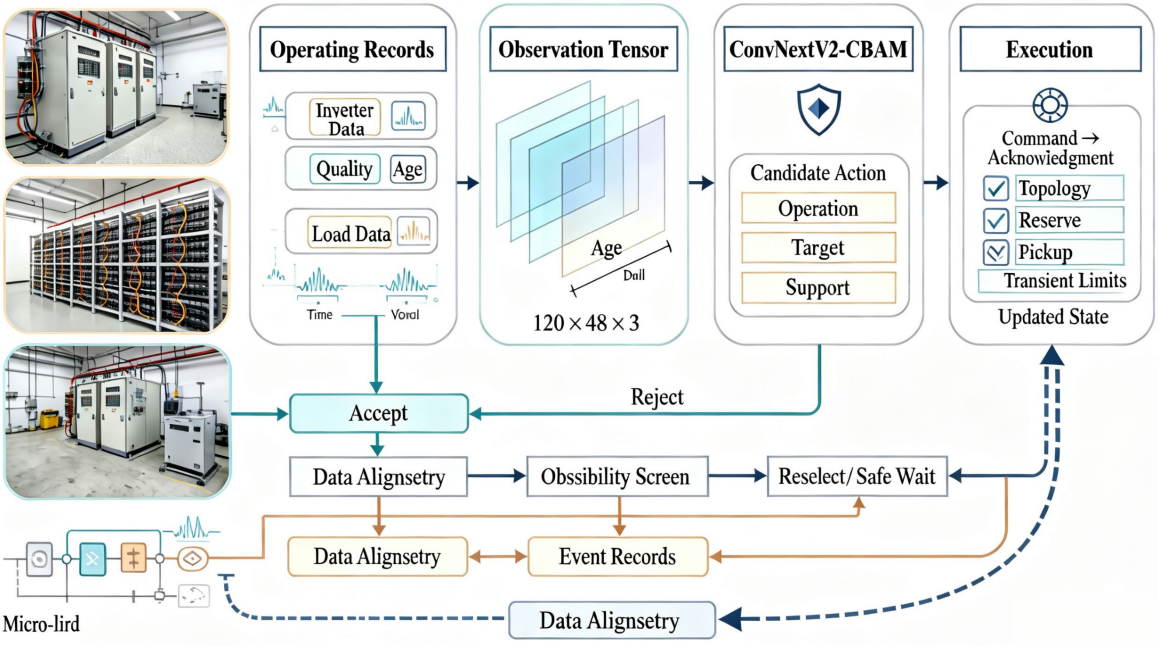



**Figure 1 Data alignment and constraint screened microgrid black start workflow**

The complete operating process is illustrated in Figure 1, where episode records feed the observation builder, the generator proposes one action, the independent screen either authorizes it or invokes recovery, and the new state returns to the record store. Execution of each authorized proposal is followed by construction of a new observation. The action semantics preserve a closed supervisory loop in which each continuation uses the acknowledged outcome of the preceding action.

### 3.2. Time Aligned Inverter and Load Inputs

The observation builder aligns signals to the supervisory decision timestamp using samples available by that instant. An inverter group contains voltage magnitude, frequency, active power, reactive power, current magnitude, current limiter status, storage state of charge, and availability. A load group contains requested active and reactive demand, priority class, interruption duration, previous pickup status, and the estimated recovery multiplier. Network descriptors identify energized buses, eligible switches, and the connectivity of each candidate target. Binary quality indicators accompany the continuous channels, and an age channel records the delay since the last valid observation, preserving measurement reliability as an explicit input. A zero filled channel with a missingness flag is processed differently from an observed zero value. The same convention applies to unavailable resources and loads pending energization.

Continuous variables are scaled using statistics computed from the training partition alone. Centering is applied to physically suitable continuous channels; binary statuses and categorical identifiers retain separate encodings. Standardization follows the equation below, whose positive denominator safeguard stabilizes the transformation for nearly constant training channels:

$$X_{k,t,c} = \frac{s_{k,t,c} - \mu_c}{\sigma_c + \varepsilon} \tag{2}$$

The index $k$ denotes a decision, $t$ a historical sample, and $c$ a continuous feature channel. The quantities $\mu_c$ and $\sigma_c$ are the training mean and standard deviation of channel $c$, while $\varepsilon$ is a small positive constant. Clipping limits are fixed from the training distribution and recorded with the preprocessing configuration. The scale statistics remain fixed throughout testing. In a deployment implementation, drift monitoring would determine whether a separately validated update is needed. The proposed electrical input has 120 historical samples and 48 ordered feature positions, with three planes carrying standardized values, quality indicators, and observation age. The feature positions represent registered measurement and status channels grouped by equipment; they are distinct from the number of physical devices. The tensor axes therefore correspond to historical time, ordered feature position, and information plane. The history spans six seconds at a twenty hertz feature rate. This rate describes the supervisory feature stream, while switching transient acquisition requires a separate measurement bandwidth. Fast peak current and protection events require independent acquisition with bandwidth suited to the transient, followed by causal aggregation into the supervisory history. Assessing peak accuracy additionally requires the acquisition rate, sensor bandwidth, and event timestamp resolution; the twenty hertz feature rate alone characterizes the supervisory input. Device positions are ordered by fixed equipment groups within a topology registry, followed by load descriptors and network status. Unused positions are padded and masked. All compared encoders receive the same ordering, and the test protocol includes device permutation checks to identify inappropriate dependence on arbitrary registry order. Extension to larger networks requires an explicit encoding for variable device counts and topology identifiers.

A causal replay check is applied when constructing each decision record. The builder fixes the decision timestamp and selects measurements whose acquisition and communication timestamps precede that instant. A late sample remains available for a later decision, with its observation age updated accordingly. Load interruption duration is measured from the last confirmed disconnection, and the corresponding acknowledgment authorizes each pickup status update. These conventions keep the feature history consistent with the information available to an online supervisor. For a device with several measurement rates, the feature extractor retains interval maxima for current and limiter activity alongside averaged operating quantities. An action identifier is stored separately from the measurement tensor so that the evaluator can reconstruct which transition followed the observation. The registry also records rated power, per unit bases, control mode eligibility, and load priority. During replay, a consistency check compares the registry version with the episode topology before features are assembled. An incompatible mapping suspends construction of that record until the device associations are resolved. Observation age continues to advance during candidate search, so eligibility must be reassessed when the frozen state exceeds the validity interval required by the electrical evaluator. The evaluator and communication interface must specify that interval before deployment.

### 3.3. Feasibility Interface and Evaluation Boundaries

The independent screen receives a proposed action, a timestamped state estimate, and the relevant device limits. Its first checks concern topology, interlocks, reference availability, and valid command syntax. The next checks concern active and reactive reserve, estimated pickup demand, storage reserve, and the operating region predicted for the energized island. A calibrated dynamic predictor or validated simulation model would then estimate transient response

over a declared acceptance interval. A compact subset of the operating envelope is written in the following equation; equipment specific interlocks and protection coordination remain additional requirements:

$$V_{min} \geq V_{lim}, max|f - f_0| \leq \Delta f_{lim}, I_{max} \leq I_{lim}, E_{res} \geq E_{lim} \tag{3}$$

Voltage and current are represented on declared per unit bases, frequency deviation in hertz, and reserve energy in kilowatt hours. The acceptance criteria use a lower voltage bound of 0.90 per unit, an absolute frequency deviation bound of 0.80 hertz, a current bound of 1.20 per unit, and a reserve floor of 120 kilowatt hours. These thresholds define the evaluation envelope and must be satisfied jointly over the assessed interval. Their application is conditional on valid measurements and the additional interlocks and protection requirements specified by the screen. A real installation would require equipment specific envelopes, permitted transient durations, protection settings, and recovery requirements. The screen also requires continuity of service for previously energized critical blocks throughout the assessed interval. Inadequate measurement quality for a required constraint triggers candidate rejection or deferral.

## 4. ConvNeXtV2 CBAM Strategy Generator

### 4.1. Encoder Architecture and Attention Adaptation

The proposed encoder uses a small ConvNeXtV2 style hierarchy to convert the aligned observation into a restoration state embedding. The stem operates on the three input planes using a patch size of two along the temporal dimension and one along the feature dimension. Subsequent stages increase channel width while aggregating temporal and grouped feature context. Four stage widths of 40, 80, 160, and 320 are proposed, with block depths of 2, 2, 6, and 2. These values specify the proposed restoration encoder configuration. Each block retains depthwise convolution, normalization, channel expansion, nonlinear transformation, global response normalization, and a residual path. Downsampling along the short feature axis preserves the representation of small device groups through the later stages. A masked pooling operation restricts the final embedding to valid observation positions.

Global response normalization compares each channel response with the channel mean of the aggregated response magnitude. For each observation, the Euclidean response magnitude is calculated over the encoded temporal and feature positions of each channel. The denominator averages those magnitudes across channels within the same observation and includes a small positive safeguard. Batch members retain separate response statistics. Within each block, the residual calibration is governed by the equation:

$$Y_c = X_c + \gamma_c X_c \frac{G_c}{mean_d(G_d) + \varepsilon} + \beta_c \tag{4}$$

Here, $G_c$ is the aggregated response magnitude of channel $c$, while $\gamma_c$ and $\beta_c$ are trainable scale and offset parameters. Both parameters are initialized to zero, so the transformation initially preserves $X_c$. This initialization maintains the residual representation while the calibration parameters are optimized. The attention module is inserted after the second and fourth encoder stages. Its channel branch uses average and maximum pooled summaries, a shared bottleneck, and a sigmoid gate. Its spatial branch uses channel pooled summaries and a convolution over the remaining temporal and feature positions. Applying channel refinement before spatial refinement gives the two transformations collected in the next equation:

$$F' = M_c(F) \odot F, F'' = M_s(F') \odot F' \tag{5}$$

The operator $\odot$ denotes elementwise multiplication, with broadcasting applied to the appropriate dimensions. Channel gates can weight learned responses associated with reserve, demand recovery, or voltage history, while the spatial gate can emphasize a transient interval within the encoded record. Such feature-weighting interpretations require channel interventions and sensitivity tests; an attention map alone describes an internal response. Padding masks must accompany the encoded history through pooling and attention inspection so that activation at unavailable positions can be distinguished from responses to valid measurements. These requirements define checks for interpreting the representation rather than additional experimentally established effects. Figure 2 presents the encoder stages, both attention insertions, the pooled state embedding, candidate decoder, and independent screening branch. The diagram separates the electrical screen from the learned network and routes rejection feedback to candidate selection.

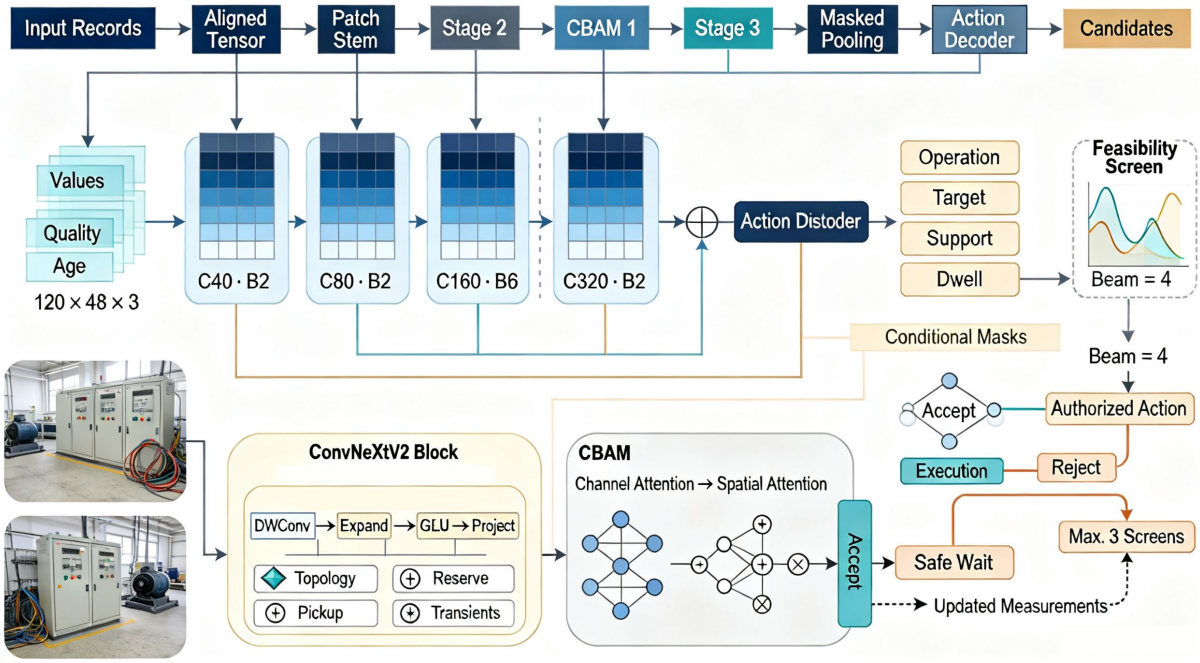



**Figure 2 ConvNeXtV2 CBAM encoder and autoregressive action decoder with an external feasibility screen.**

## 4.2. Candidate Decoding and Training Specification

The pooled state embedding is combined with a compact encoding of the accepted action history. An autoregressive decoder emits operation type, target device, optional support level, and dwell category in a fixed order. Conditional masks restrict candidate scoring to syntactically valid targets. A pickup operation targets an unserved eligible load, while reference enabling targets an available device with grid forming capability. These masks enforce command syntax and state compatibility before electrical screening. Candidate ranking uses the decoder score together with a learned estimate of subsequent unserved energy, while the independent screen determines whether the ranked action may be executed. The candidate beam width is four, while at most three ranked, complete action candidates undergo electrical screening at each supervisory decision. These limits govern different operations: beam width bounds decoding search, and the screening budget bounds calls to the electrical evaluator. A wait action remains available whenever its data quality and reserve conditions are satisfied.

Training uses expert or optimization derived action records associated with a specified electrical model. Every target action must be associated with the preaction observation, the verified action outcome, and the version of the model that accepted it. Action likelihood, a value target, and an auxiliary violation label contribute to the training objective through the weighted equation below:

$$L = L_{act} + \alpha L_{val} + \beta L_{viol} \tag{6}$$

The action term $L_{act}$ is cross entropy over valid tokens, the value term $L_{val}$ is a normalized regression loss on future weighted unserved energy, and the violation term $L_{viol}$ is binary cross entropy for the outcome of the independent evaluator. The violation term trains a diagnostic head, while authorization remains with the independent evaluator. Coefficients of 1.0,0.2, and 0.5 are used for the respective contributions after their individual scales are inspected. The action loss has unit weight, the value loss has weight 0.2, and the violation loss has weight 0.5. These coefficients specify the reported training configuration; their sensitivity remains a separate evaluation question. The auxiliary violation output supports diagnosis, while the independent electrical evaluator retains authorization authority. Architecture comparisons require disclosure of optimizer settings, stopping criteria, random seeds, data augmentation, and computational resources. The training configuration uses AdamW with a learning rate of 0.0003, weight decay of 0.05, batch size of 64, and an upper limit of 150 epochs. Validation loss and screened sequence behavior jointly determine the selected checkpoint. Temporal masking and sensor dropout should follow physically defensible ranges and retain the original observation quality indicators. Published image weights should be treated as an optional initialization experiment because their preprocessing and learned feature semantics differ from those of restoration data. Any claimed benefit of pretraining would require an otherwise matched comparison against random initialization.

When the expert evaluator identifies several admissible actions with comparable restoration cost, the training record should retain their eligibility and cost information. The selection or weighting of those alternatives must be declared with the target-generation procedure. This requirement describes how ambiguous supervision should be documented; the aggregate comparison does not isolate the effect of alternative-target weighting. The decoder can represent this

ambiguity through a distribution over eligible operation tokens, followed by conditional target and dwell predictions. Validation should then examine both sequence cost and the fraction of proposals accepted by the independent screen. For rejected training candidates, the stored violation label should be conditioned jointly on the evaluated state and proposed action. This distinction allows two different proposed pickups from the same observation to receive different labels. Value targets require a common horizon and fixed priority weights so that their scale remains comparable across episodes. Checkpoint selection should use those fixed definitions together with the declared validation partition. Once selected, the encoder statistics, decoder weights, token registry, and screening interface version are frozen as one reproducible controller configuration.

### 4.3. Online Screening and Recovery Logic

At each supervisory decision, the controller fixes a causally complete observation and submits it to the decoder to obtain a ranked candidate list. The first candidate is checked for valid identifiers and allowable state transitions. The screen then evaluates resource reserve, pickup uncertainty, and the required electrical envelope using the current state estimate. Candidate authorization requires complete evaluation and successful satisfaction of every applicable check. Authorization concerns an individual proposed action; episode acceptance additionally concerns the restoration outcome under the evaluation envelope and priority-service target. If it is rejected, the reason is stored using a fixed vocabulary such as insufficient reserve, predicted overcurrent, reference unavailable, stale measurements, or incompatible topology. The next candidate may then be considered within a strict perdecision budget. Repeated rejection exhausts a fixed search budget, after which the supervisor obtains a fresh observation before further candidate assessment. When the budget expires, the controller requests a wait or invokes a deterministic recovery policy only while its reserve, observation quality, and operating envelope remain admissible. If those conditions fail, the supervisor must defer to the installed protection and validated recovery procedures. Waiting itself consumes restoration time and can reduce the available energy reserve.

Execution proceeds through an acknowledged command channel. Command transmission, switch acknowledgment, and confirmed load service are recorded as distinct events. The observation builder updates restoration status after the corresponding command acknowledgment, confirmation of physical load service, and satisfaction of the required settling condition. Until these conditions are verified, the transition remains pending in the action history. A transmitted command and an acknowledged switch transition therefore remain separate from confirmed service delivery. During the dwell interval, fast protection and local regulation retain authority. A protection operation invalidates the pending continuation and causes the supervisor to acquire a new state. During partial pickup failure, verified acknowledgments preserve the consistency of the action history and available reserve estimate. The recorded transition includes the proposal, all rejections considered before authorization, the issued command, the acknowledged physical state, and any protection event. Separate fields support diagnosis and preserve the distinction between issued commands and acknowledged physical transitions in training records.

Computational evaluation must distinguish encoder time, candidate decoding time, electrical screening time, and communication delay. Complete controller latency includes encoder computation, candidate decoding, dynamic screening, and communication. Network inference, electrical screening, and communication are evaluated as separate contributions to the supervisory decision interval. Confidence is monitored through validation calibration and rejected proposal frequency; a low confidence state can request deferral, while hard operating restrictions remain in force. Device failures and topology changes outside the declared configuration trigger the validated fallback. The recovery policy must be validated independently, including its behavior under energy depletion and prolonged missing observations. This design specifies how the learned component can participate in black start supervision while preserving explicit operational responsibilities. The data interface, electrical evaluator, command acknowledgment, and recovery policy must use compatible device identifiers and synchronized operating states throughout a restoration episode.

## 5. Performance Evaluation and Engineering Analysis

### 5.1. Evaluation Records and Transient Diagnostics

The evaluation examines operating coverage, transient response, restoration efficiency, and the supervisory requirements of the proposed strategy generator. Scenario based restoration studies motivate joint variation of resource availability and pickup demand when defining an evaluation domain [25]. The collection contains 5100 operating records, partitioned into 3600 training, 900 validation, and 600 testing identifiers at the episode level. These counts define the partitioning units; historical windows derived from an episode must remain within its assigned partition. The number of training windows is distinct from the number of episode identifiers. Episode level separation addresses the dependence introduced when overlapping histories from one restoration trajectory enter different partitions [26]. Initial storage charge spans 0.35 to 0.85, the recovery multiplier spans 1.0 to 2.2, and five feeder stress classes are combined with missingness levels of 0, 5, 10, 15, and 20 percent. Table 1 specifies a 400-volt, 50 hertz island containing three inverter resources, eight load groups, 750 kilovolt amperes of aggregate inverter capacity, 500 kilowatts of requested

demand, and 1200 kilowatt hours of storage. The critical and essential service target is 380 kilowatts. This priority-service target is distinct from the 500-kilowatt demand base used to normalize terminal service; full restoration refers to all requested demand. A common 600-episode denominator supports comparisons of acceptance and terminal service, whereas accepted episode duration uses the accepted subset of each profile. The records retain resource conditions, load recovery descriptors, acceptance flags, and pickup response histories. Their common identifiers connect terminal performance with the sequence of supervisory decisions that produced it. Reserve and recovery demand are evaluated jointly, allowing low charge conditions to coincide with demanding pickup requirements. This combination challenges the decision mechanism because both transient support and sustained energy availability influence the next admissible action. Observation loss is recorded separately from electrical stress so that the assessment can distinguish a difficult operating condition from an inadequately observed state.

**Table 1. Evaluation specification and controller settings**

| Item | Value | Interpretation |
|---|---|---|
| Nominal island | 400 V and 50 Hz | Electrical base |
| Resources and demand | 3 inverters and 8 load groups | 750 kVA and 500 kW totals |
| Storage | 1200 kWh | Installed storage energy |
| Scenario records | 3600 train 900 validation 600 test | Episode partition |
| Observation tensor | 120 by 48 by 3 | Values quality and age planes |
| Feature stream | 20 Hz over 6 s | Supervisory features only |
| Acceptance envelope | 0.90 pu 0.80 Hz 1.20 pu | Voltage floor frequency deviation current ceiling |
| Candidate budget | Beam 4 and at most 3 screens | Supervisory setting |

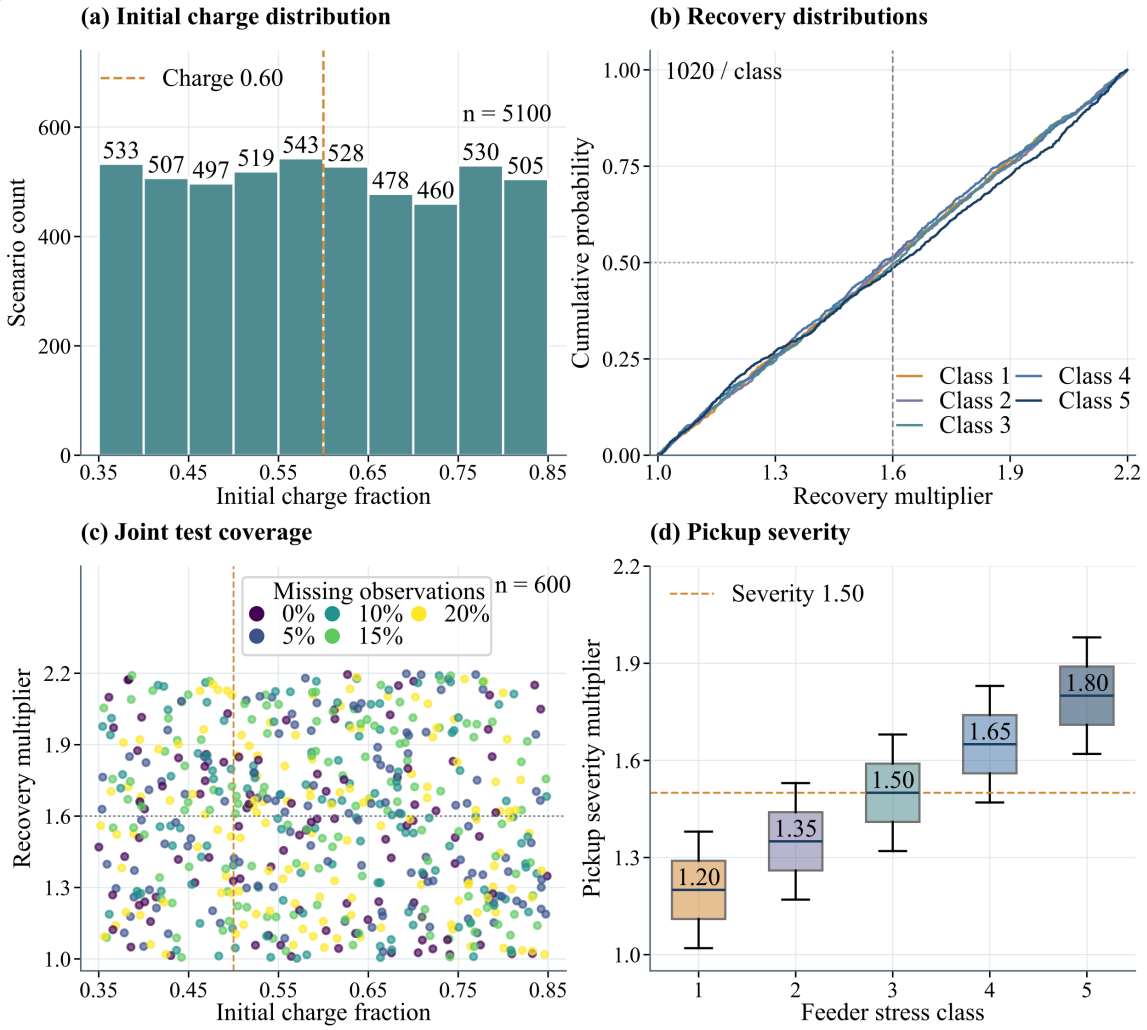



**Figure 3 Operating condition coverage and stress composition. (a) Storage charge histogram. (b) Recovery multiplier empirical cumulative distributions. (c) Joint charge and recovery coverage with missingness levels. (d) Pickup severity by stress class**

Figure 3(a) presents the storage charge distribution across 5100 descriptors, with ten bins spanning 0.35 to 0.85 and a reference at 0.60. Counts vary from 460 to 543, so the most populated interval contains approximately 18.0 percent more records than the least populated interval. The count variation provides coverage throughout the reserve range while retaining enough observations in each interval to inspect differences in initial operating conditions. Figure 3(b) presents recovery multiplier cumulative distributions for five classes containing 1020 records each. All curves occupy the common 1.0 to 2.2 interval, and their proximity around the 1.6 reference shows that stress labels are separated from the sampled recovery marginal. Evaluation should consequently assess the joint use of recovery observations and stress descriptors in the generated decisions.

Figure 3(c) presents the 600 test descriptors jointly by charge, recovery multiplier, and five missingness levels. Points on both sides of the 0.50 charge reference and the 1.60 recovery reference expose four reserve and pickup combinations. The low charge and high multiplier region are especially relevant to reserve screening because limited initial energy coincides with elevated requested pickup. Figure 3(d) presents pickup severity distributions with medians of 1.20, 1.35, 1.50, 1.65, and 1.80 across increasing stress classes. Each distribution spans 0.36, while adjacent medians differ by 0.15. The resulting overlap makes intermediate cases more informative than a classification problem with disjoint severity intervals. Recovery multiplier and pickup severity describe different aspects of restoration demand: the former characterizes recovered consumption, whereas the latter characterizes the intensity of the energization event.

Inverter energization studies connect current limiting with the ability to sustain a reference during transient demand [27]. Figure 4(a) presents six voltage response trajectories with minima of 0.875, 0.887, 0.899, 0.911, 0.923, and 0.935 per unit. Relative to the 0.90 floor, the first three trajectories have deficits of 0.025, 0.013, and 0.001 per unit, whereas the remaining trajectories retain positive margins of 0.011, 0.023, and 0.035 per unit. The third trajectory is therefore close to the decision boundary, making the treatment of measurement accuracy consequential when this diagnostic is transferred to equipment records. Figure 4(b) presents corresponding current peaks of 1.34, 1.30, 1.26, 1.22, 1.18, and 1.14 per unit against a 1.20 ceiling. The fourth profile retains a positive voltage margin while exceeding the current ceiling by 0.02 per unit. Its different classification under the two indicators demonstrates why an acceptance decision requires simultaneous envelope checks. The fifth and sixth profiles retain current margins of 0.02 and 0.06 per unit, respectively. Both panels mark pickup at two seconds and highlight the following transient interval. The response curves follow similar recovery stages, enabling comparison of the voltage depression and current surge associated with the pickup event. Their later convergence toward the common operating level demonstrates the value of retaining transient extrema alongside terminal measurements. For future dynamic records, the duration of a threshold excursion, the sampling rate, and the protection response should accompany peak values. These additional quantities determine whether an apparent crossing corresponds to a permitted short transient or an unacceptable operating event.

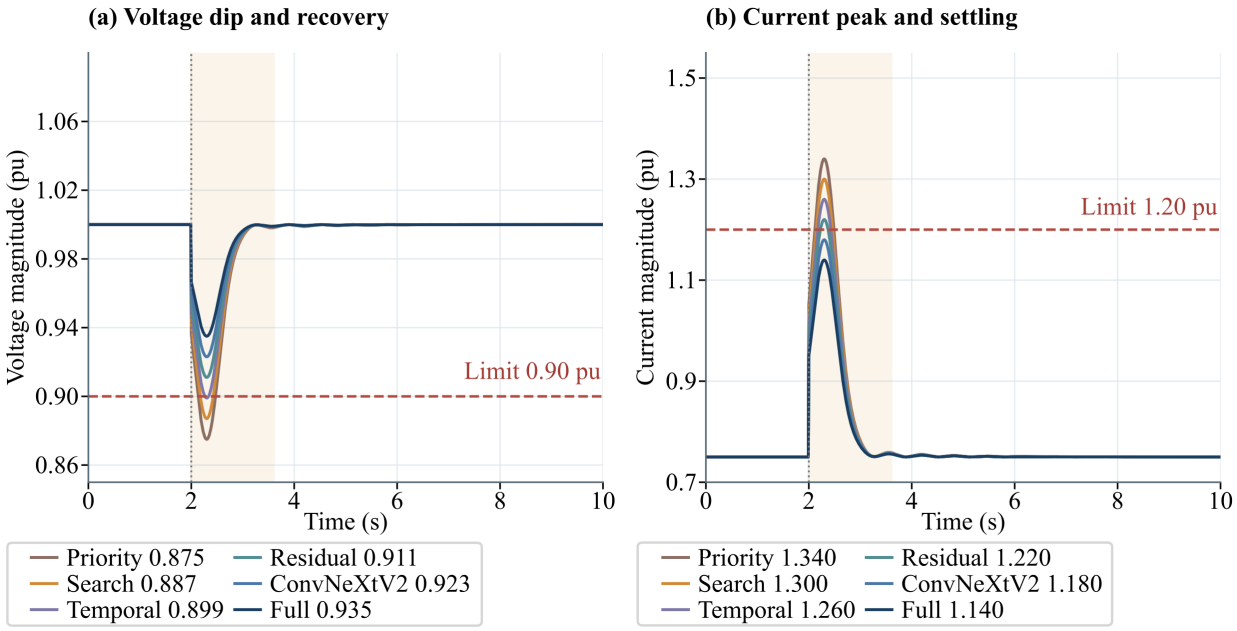



**Figure 4 Pickup voltage and current responses. (a) Voltage magnitude with the 0.90 pu floor. (b) Current magnitude with the 1.20 pu ceiling**

Electrical transient validation requires agreement in event timing and limiting behavior as well as waveform amplitude [28]. The response trajectories support checks of extrema extraction, timestamp alignment, and the direction of operating margins. The common event window separates the initial pickup disturbance from the subsequent recovery interval, while minima and maxima identify the closest approach to each operating limit. Paired controller evaluation

likewise relies on common episode conditions to reduce differences arising from the test collection itself [29]. The six methods use identical testing identifiers, common operating limits, and consistent episode termination rules. Acceptance counts are 474, 504, 528, 546, 558, and 576, while accepted duration means are 142, 128, 118, 110, 106, and 98 seconds. The accepted duration intervals extend 28, 26, 24, 22, 20, and 18 seconds on either side of the respective means. These differences in spread distinguish methods with similar central durations but different exposure to prolonged restoration episodes. Voltage, current, frequency deviation, reserve, and critical service fields accompany every record, enabling each acceptance flag to be checked against the applicable operating envelope. Failed records retain terminal service and their overcurrent indicator, while their accepted duration field remains empty. This convention separates incomplete attempts from completed restoration episodes and preserves a consistent denominator for service assessment. Each performance metric retains its episode identifier and applicable denominator, allowing differences in acceptance to be interpreted alongside differences in completed service. Duration distributions are examined together with terminal service because the accepted subset can change when the decision mechanism admits a wider range of operating conditions. The comparison of accepted-subset means is descriptive and combines speed differences with possible changes in the accepted population. A paired duration analysis would require the common subset accepted by both methods, while uncertainty estimates would require the episode-level outcomes.

### 5.2. Comparative Performance and Component Diagnostics

Fair control benchmarks align information access, action eligibility, and episode termination rules before interpreting differences between methods [30]. Table 2 applies a common reporting structure to the six methods. Figure 5(a) presents accepted counts of 474, 504, 528, 546, 558, and 576, corresponding to 79, 84, 88, 91, 93, and 96 percent of the 600 episodes. The selected 95 percent reference corresponds to 570 accepted episodes. Full exceeds that reference by six episodes, whereas ConvNeXtV2 remains twelve episodes below it. Their eighteen-episode separation is three percentage points, illustrating why a percentage difference should also be expressed as a count.

**Table 2. Comparative performance on 600 common testing episodes**

| Method | Accepted episodes | Mean accepted duration s | Mean terminal service percent | Inference delay ms |
|---|---|---|---|---|
| Priority | 474 of 600 | 142.0 | 83.0 | 1.0 |
| Search | 504 of 600 | 128.0 | 86.0 | 15.0 |
| Temporal | 528 of 600 | 118.0 | 89.0 | 4.5 |
| Residual | 546 of 600 | 110.0 | 91.0 | 5.8 |
| ConvNeXtV2 | 558 of 600 | 106.0 | 92.0 | 6.8 |
| Full | 576 of 600 | 98.0 | 95.0 | 7.6 |

Figure 5(b) pairs the mean accepted durations of 142, 128, 118, 110, 106, and 98 seconds with ninety fifth percentiles of 167.2, 151.4, 139.6, 129.8, 124.0, and 114.2 seconds. The tail excess above the mean contracts from 25.2 seconds for Priority to 16.2 seconds for Full. The Full and ConvNeXtV2 profiles differ by eight seconds at the mean and 9.8 seconds at the upper percentile. Across the duration distributions, the shift in the upper tail is therefore slightly larger than the shift at the center. Acceptance and duration answer separate operational questions: the former counts episodes meeting the envelope and service target, while the latter describes completion time within the corresponding accepted subset. The joint comparison connects the frequency of successful completion with the time required by accepted episodes. Attention effects are assessed through matched network configurations and a common decision budget.

Restoration metrics distinguish recovery speed from the quantity and priority of service delivered during an interruption [31]. Figure 6(a) presents six accepted duration distributions spanning 114 to 170, 102 to 154, 94 to 142, 88 to 132, 86 to 126, and 80 to 116 seconds. Each violin retains its own accepted sample count, making the different population sizes explicit. The overlapping intervals show that separated profile means can coexist with substantial overlap between individual episodes. The 100 second reference intersects the final four distributions, although their means range from 98 to 118 seconds. Crossing that reference in an individual episode therefore provides a different classification from placing a profile mean below it. Figure 6(b) presents network inference delays of 1.0, 15.0, 4.5, 5.8, 6.8, and 7.6 milliseconds against the acceptance percentages, with mean terminal service encoded as 83, 86, 89, 91, 92, and 95 percent. The Full and ConvNeXtV2 points differ by 0.8 milliseconds of inference delay, three percentage points of acceptance, and three percentage points of mean service. Computational benchmarking requires a stated processor, synchronization procedure, batch size, and warmup policy when delays are measured [32]. The horizontal coordinates represent network inference, separating model computation from the remaining supervisory operations. The 10-millisecond reference separates the Search inference delay from the other five values; assessment against the complete supervisory deadline also requires electrical screening and communication delays. The reported inference delays characterize the model-computation component. Their transfer to another platform requires the processor, numerical

precision, batch size, synchronization procedure, and warmup policy to be specified. Complete deadline assessment additionally requires measured screening and communication delays under that platform configuration.

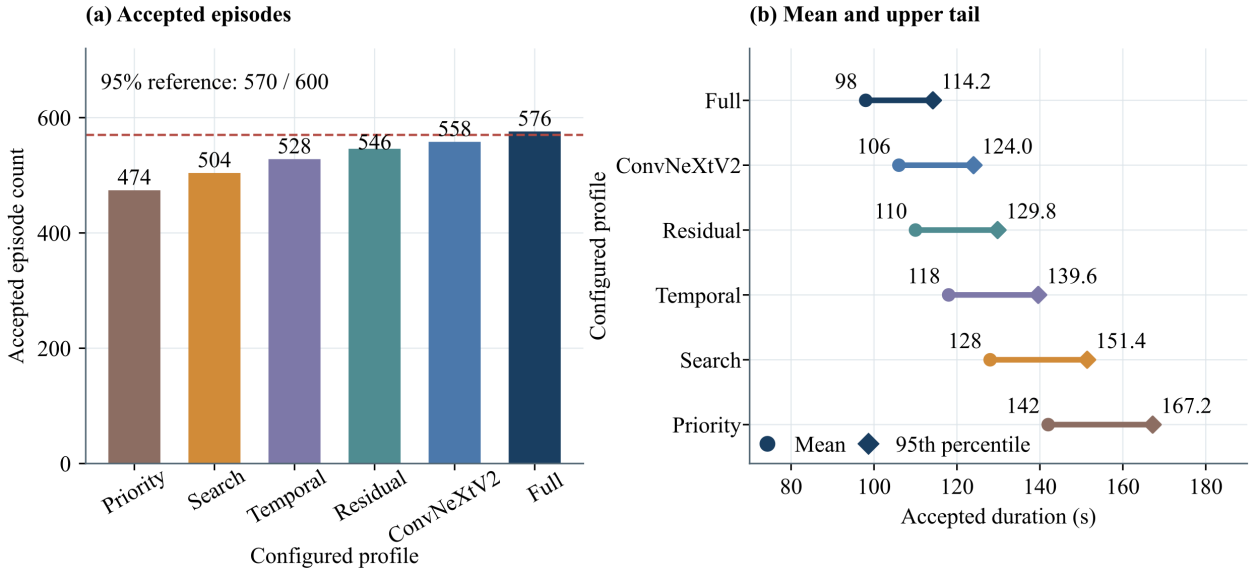



**Figure 5 Acceptance and restoration duration. (a) Accepted counts with a 570 of 600 reference. (b) Mean and ninety fifth percentile duration among accepted episodes**

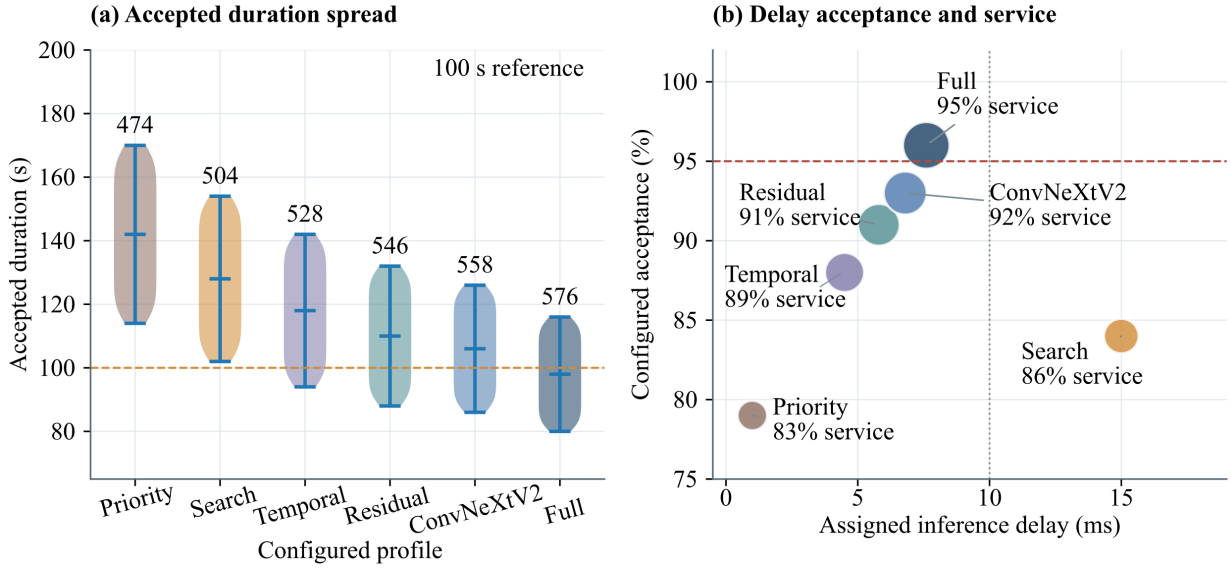



**Figure 6 Duration distribution and computational tradeoffs. (a) Accepted duration violin distributions. (b) Network inference delay versus acceptance with terminal service encoded by marker area**

Controlled component studies hold observations, optimization effort, and candidate budgets constant while changing the mechanism under examination [33]. For the proposed architecture, the relevant mechanisms are feature normalization, attention refinement, action decoding, and independent screening. Full and ConvNeXtV2 differ by eighteen accepted episodes and eight seconds of mean accepted duration. This configuration-level comparison describes the combined outcome under the reported protocol. Attribution to attention or normalization requires separately controlled component results and variability across training seeds. Component diagnostics distinguish improvement in proposal quality from the contribution of execution filtering by recording the candidate rank and screening outcome at each decision. If attention improves representation, a useful diagnostic is the fraction of first proposals that pass the independent screen under the same scenario. If screening accounts for most of the improvement, the intervention count may increase even when final acceptance rises. These two mechanisms impose different computational demands on candidate search and electrical evaluation. Increasing the candidate budget can raise acceptance while increasing decoding and screening delay. A matched comparison should consequently retain beam width four and the limit of three screens, and record which candidate is eventually authorized. Channel removal can test dependence on voltage, current limiting status, and reserve descriptors, while a permutation test can expose sensitivity to arbitrary device ordering. Normalization statistics should remain fixed during these interventions so that input scaling and architecture effects remain distinguishable. For duration comparisons, the accepted subset should be accompanied by service over all

episodes, since a change in acceptance can alter which scenarios contribute to the reported mean. The training protocol should log each random seed and checkpoint selection rule, permitting the observed separation to be compared with variability arising from parameter initialization.

Intervention frequency is an operational quantity in constrained policy evaluation because rejected proposals consume decision effort and can extend service interruption [34]. The failure counts are 126, 96, 72, 54, 42, and 24, which complement the accepted counts on the common denominator. The reduction from 42 to 24 failures represents eighteen fewer failures, or 42.9 percent of the ConvNeXtV2 failure count, while the acceptance difference remains three percentage points. The relative failure reduction uses baseline failures as its denominator, whereas the acceptance difference uses all test episodes. These failure counts describe episode outcomes. Rejected proposals are action-level events and require separate logs; their frequency cannot be inferred from the difference in failed episodes. Mean terminal service rises from 92 to 95 percent between the same profiles, corresponding to 460 and 475 kilowatts on the 500-kilowatt demand base. The fifteen-kilowatt difference describes average end of horizon service across every record, including incomplete attempts. The failure records include the overcurrent indicator associated with the interrupted restoration attempt. An electrical experiment would additionally distinguish loss of reference, undervoltage, frequency excursion, reserve depletion, synchronization failure, and acknowledgment timeout. Rejection of a proposed command should be logged separately from violation after an authorized command, since the latter also indicates a possible error in the electrical evaluator or state estimate. Waiting preserves a feasible operating condition while accumulating further unserved energy, so a successful rejection can still carry a restoration cost. For episodes with several rejected candidates, recording the rejection order and the age of the frozen observation would expose whether candidate search continues beyond the validity of the state used by the screen.

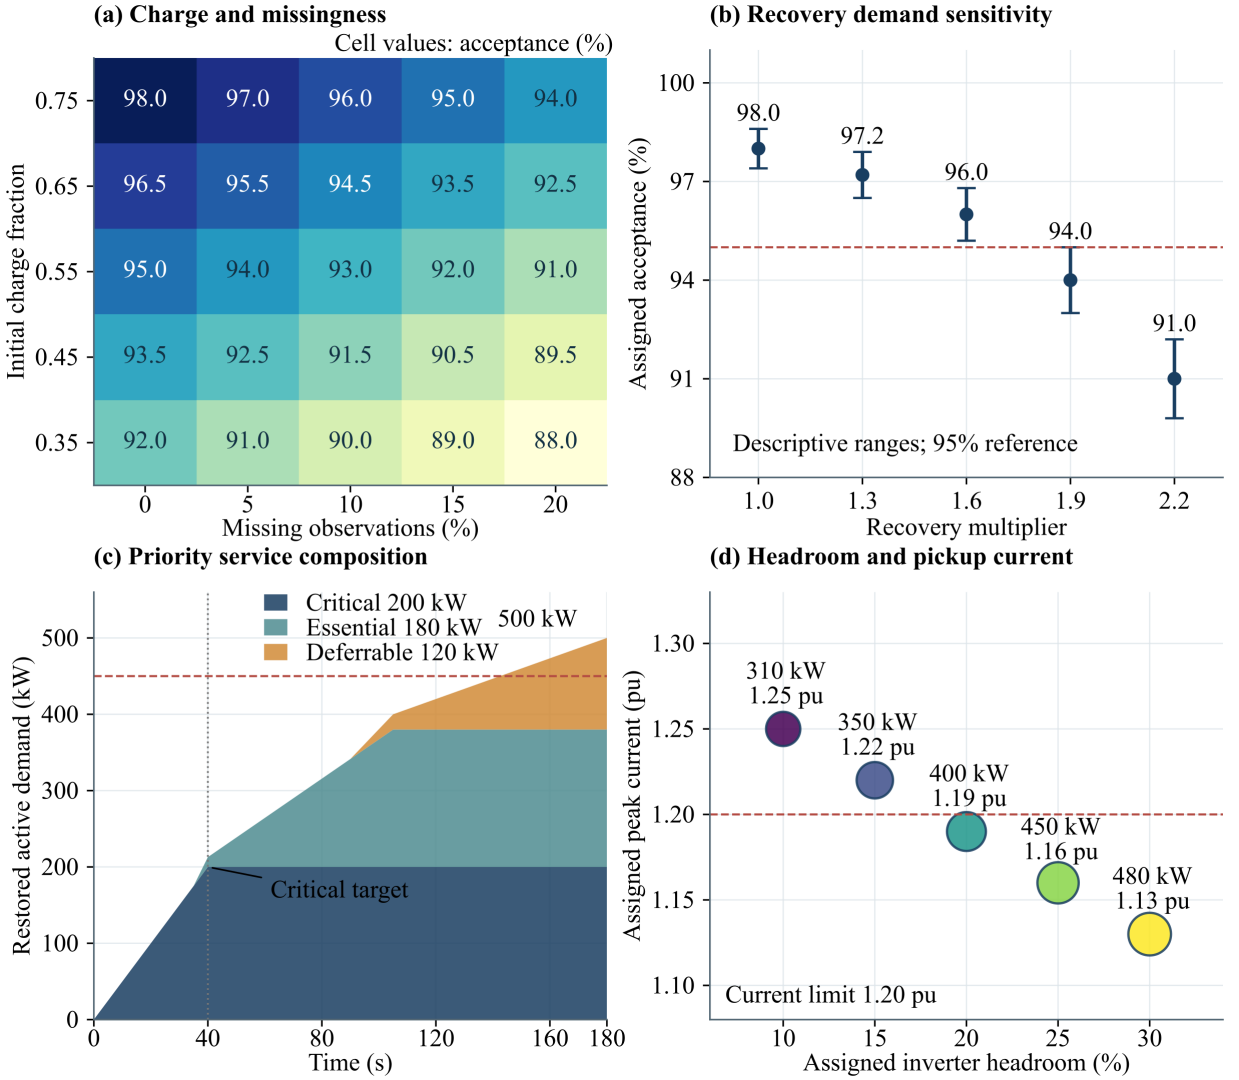



**Figure 7 Sensitivity and service composition. (a) Acceptance over charge and missingness. (b) Acceptance versus recovery multiplier with descriptive ranges. (c) Service by priority class. (d) Headroom current and restored demand**

### 5.3. Sensitivity and Operational Sequencing

Robustness evaluation in learned power system control examines coupled operating and observation changes because their interaction can alter action eligibility [35]. Figure 7(a) presents Full acceptance over charge levels of 0.35, 0.45, 0.55, 0.65, and 0.75 and missingness levels from 0 to 20 percent. Cell values span 88.0 to 98.0 percent. At each fixed charge level, increasing missingness across the displayed range reduces acceptance by four percentage points; increasing charge across its displayed range raises it by six percentage points. Across the evaluated operating conditions, the 95 percent reference is met or exceeded at charge 0.55 with complete observations, at charge 0.65 with five percent missingness, and at charge 0.75 with fifteen percent missingness. These contour crossings express an explicit relationship between initial resource reserve and observation availability. Figure 7(b) presents acceptance values of 98.0, 97.2, 96.0, 94.0, and 91.0 percent for recovery multipliers of 1.0, 1.3, 1.6, 1.9, and 2.2. Successive declines are 0.8, 1.2, 2.0, and 3.0 percentage points, so the assumed degradation becomes steeper at higher recovery demand. The displayed ranges have widths of 0.6, 0.7, 0.8, 1.0, and 1.2 percentage points. They are treated here as descriptive plotting intervals; a confidence interpretation requires the construction rule, sample counts, and repeated outcomes for each operating condition. The value at multiplier 1.9 lies one point below the reference, while its upper range endpoint touches that reference. This boundary case requires consideration of both range endpoints because the central value and the upper endpoint imply different positions relative to the acceptance reference. Device specific missingness and independently generated pickup responses would be necessary to identify whether the same pattern persists in an implemented controller.

Resilience assessment uses service trajectories to account for the timing and importance of recovered demand [36]. Figure 7(c) presents critical, essential, and deferrable components with targets of 200, 180, and 120 kilowatts. The critical component reaches its target at 40 seconds, the essential component at 105 seconds, and the deferrable component at 180 seconds. At 40 seconds, the critical fraction is complete while essential service is approximately 12.9 kilowatts and deferrable service remains zero. At 105 seconds, critical and essential service total 380 kilowatts, while deferrable service contributes 20 kilowatts, giving 400 kilowatts overall. The component trajectories identify the priority composition underlying the aggregate service increase. Figure 7(d) presents headroom levels of 10, 15, 20, 25, and 30 percent paired with peak currents of 1.25, 1.22, 1.19, 1.16, and 1.13 per unit and restored demand of 310, 350, 400, 450, and 480 kilowatts. The current ceiling is crossed between the second and third operating points. At twenty percent headroom, the current margin is 0.01 per unit; the final point retains 0.07 per unit and serves eighty additional kilowatts. The combination of current coordinates and marker areas therefore distinguishes electrical margin from service benefit. The joint variation of headroom, peak current, and restored demand identifies the operating margin associated with an additional pickup. Applying the same diagnostic to equipment requires apparent power reserve, reactive support, voltage, and limiter status to be expressed on compatible bases, because identical active power headroom can accompany different transient current capability.

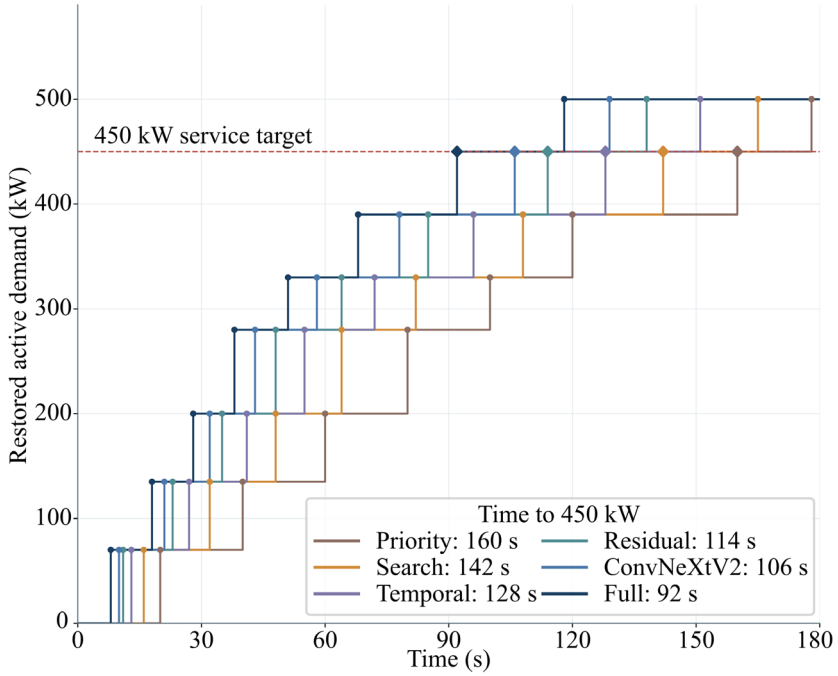



**Figure 8 Staged service trajectories**

Staged restoration studies coordinate pickup intervals with the changing support capability of the energized island [37]. Figure 8 presents six representative sequences, each containing eight pickups of 70, 65, 65, 80, 50, 60, 60, and 50

kilowatts. The cumulative levels are therefore 70, 135, 200, 280, 330, 390, 450, and 500 kilowatts. Full reaches 450 kilowatts at 92 seconds, followed by ConvNeXtV2 at 106 seconds, Residual at 114 seconds, Temporal at 128 seconds, Search at 142 seconds, and Priority at 160 seconds. The fourteen second difference between the two encoder profiles at that threshold is larger than their eight second difference in aggregate mean accepted duration because the plotted trajectory and the aggregate record describe different evaluation quantities. Full reaches the first 200 kilowatts at 28 seconds, compared with 32 seconds for ConvNeXtV2 and 60 seconds for Priority. Its early service advantage therefore begins before the final optional blocks are picked up. The interval from 390 to 450 kilowatts is twenty-four seconds for Full and twenty-eight seconds for ConvNeXtV2; the following interval to 500 kilowatts is twenty-six and twenty-three seconds, respectively. The faster threshold crossing thus coexists with a longer final waiting interval. Examining individual plateaus locates the remaining delay within the sequence and complements the terminal completion time. Pickup acknowledgments and settling records would determine whether such a plateau reflects electrical recovery, measurement availability, or deliberate reserve management in an implemented system.

Controller hardware evaluation relates supervisory decisions to measured command timing and equipment feedback under the intended communication interface [38]. The representative Full sequence completes at 118 seconds, while the corresponding ConvNeXtV2, Residual, Temporal, Search, and Priority sequences complete at 129, 138, 151, 165, and 178 seconds. For Full, the last fifty kilowatts occupy twenty-six seconds, approximately 22.0 percent of the total sequence duration, despite representing ten percent of nominal demand. That concentration of elapsed time in the final block identifies a specific interval for inspecting reserve release and acknowledgment delay. Across the two encoder sequences, the lead decreases from fourteen seconds at the 450-kilowatt threshold to eleven seconds at full service. The final pickup therefore contributes differently to early service and terminal duration. The 118-second Full sequence is a representative trajectory, whereas the 98-second value in the comparative results is a mean over accepted episodes. These quantities have different sampling scopes and support different interpretations. In a timestamped implementation, this interval should be decomposed into observation acquisition, candidate generation, electrical evaluation, command transmission, acknowledgment, and settling. The components have distinct diagnostic roles: acknowledgment delay extends the command confirmation interval, while a settling interval allows electrical recovery. A rejected candidate between the seventh and eighth pickups would additionally create a gap between decision effort and completed switching. The event record should retain its rejected target, reason, observation age, and subsequent authorized operation. Comparing these fields with the plateau duration would allow an extended interval to be traced to the responsible stage of the supervisory process.

## 6. Conclusion

This paper developed a ConvNeXtV2 CBAM strategy generator that connects inverter histories and load restoration records to supervisory black start actions. The framework separates observation alignment, hierarchical feature encoding, candidate decoding, independent feasibility screening, and acknowledged execution. This organization makes each decision traceable to the available measurements and applicable operating constraints. The evaluation examines acceptance, restoration duration, terminal service, transient margins, and computational requirements through complementary indicators. Attention refinement and constraint screening occupy distinct roles within the controller: one shapes the representation used to propose actions, while the other determines their operational eligibility during the evolving restoration process.

Several limitations require further investigation. The feature representation depends on a consistent device registry and a bounded observation window, which can restrict transfer to networks with different resource populations or measurement rates. The electrical screen also depends on accurate pickup demand, converter limits, and timely state information. Errors in these inputs can change the assessed operating margin and increase rejection frequency. Attention responses describe feature weighting and require complementary intervention analysis for engineering interpretation. The interaction among transformer energization, unbalanced loads, protection operation, and communication delay also warrants closer examination under demanding restoration conditions and heterogeneous equipment configurations. The reported aggregates leave component attribution, run-to-run variability, acquisition bandwidth, and complete controller timing as limits on engineering interpretation.

Future research should extend evaluation to changing topology, variable inverter populations, and broader load recovery conditions. Matched component studies should isolate encoder, attention, decoder, and screening effects while retaining common observations and decision budgets. Hardware evaluation should examine command acknowledgment, adverse sensing conditions, protection interactions, and recovery behavior. Deployment assessment should first establish valid sensing intervals and equipment-specific envelopes, then measure the complete decision path on the intended controller platform. Adaptive representations may improve the treatment of observation age and changing device availability, provided their updates preserve consistent acceptance criteria. Further work should also investigate calibration drift and the coordination of restoration across multiple islands. These developments support an auditable supervisory controller whose decisions remain interpretable as operating conditions and network structure evolve.

# 7. Declarations

## 7.1 Author Contributions

Marcel Jasiński contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Władysław Gustaw Jach contributes to validation, analysis, investigation, data collection, draft preparation. Marcel Knapik contributes to conceptualization, validation, analysis. Dominik Błasik contributes to software, validation, analysis, investigation. All authors have read and agreed with the manuscript before its submission and publication.

## 7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

## 7.3 Institutional Review Board Statement

Not applicable.

## 7.4 Informed Consent Statement

Not applicable.

## 7.5 Date Availability

Not applicable.

## 7.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## 7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

# 8. References

[1] Seo, G. S., Wang, W., & Mirafzal, B. (2025). Microgrid black start challenges: the role of grid-forming inverters. IEEE Access, 13, 198087-198101. https://doi.org/10.1109/ACCESS.2025.3634530

[2] Khan, M., Wu, W., & Li, L. (2024). Grid-forming control for inverter-based resources in power systems: A review on its operation, system stability, and prospective. IET Renewable Power Generation, 18(6), 887–907. https://doi.org/10.1049/rpg2.12991

[3] Baeckeland, N., Chatterjee, D., Lu, M., Johnson, B., & Seo, G.-S. (2024). Overcurrent limiting in grid-forming inverters: A comprehensive review and discussion. IEEE Transactions on Power Electronics, 39(11), 14493–14517. https://doi.org/10.1109/TPEL.2024.3430316

[4] Xie, D., Xu, Y., Nadarajan, S., Viswanathan, V., & Gupta, A. K. (2024). Dynamic frequency-constrained load restoration considering multi-phase cold load pickup behaviors. IEEE Transactions on Power Systems, 39(1), 107–118. https://doi.org/10.1109/TPWRS.2022.3225798

[5] Hu, Z., Yang, X., Ding, L., Li, Z., He, Y., Wu, Q., & Wen, J. (2025). Measurement-correlated resilience enhancement for R-SOP-integrated distribution systems with voltage security. IEEE Transactions on Smart Grid, 16(6), 4517–4533. https://doi.org/10.1109/TSG.2025.3594669

[6] Zheng, J., Maharjan, S., & Wang, Z. (2026). Grid edge intelligence-assisted model predictive framework for black start of distribution systems with inverter-based resources. IEEE Transactions on Smart Grid, 17(3), 1863–1874. https://doi.org/10.1109/TSG.2026.3652515

[7] Pang, K., Wang, C., Hatziargyriou, N. D., & Wen, F. (2024). Dynamic restoration of active distribution networks by coordinated repair crew dispatch and cold load pickup. IEEE Transactions on Power Systems, 39(2), 4699–4713. https://doi.org/10.1109/TPWRS.2023.3309862

[8] Raghuvamsi, Y., & Teeparthi, K. (2024). Distribution system state estimation with Transformer-Bi-LSTM-based imputation model for missing measurements. Neural Computing and Applications, 36(3), 1295–1312. https://doi.org/10.1007/s00521-023-

09097-5

[9] Xu, R., Zhang, C., Zhang, D., Dong, Z. Y., & Yip, C. (2024). Adaptive robust load restoration via coordinating distribution network reconfiguration and mobile energy storage. IEEE Transactions on Smart Grid, 15(6), 5485–5499. https://doi.org/10.1109/TSG.2024.3404776

[10] Vu, L., Vu, T., Vu, T. L., & Srivastava, A. (2024). Multi-agent deep reinforcement learning for distributed load restoration. IEEE Transactions on Smart Grid, 15(2), 1749–1760. https://doi.org/10.1109/TSG.2023.3310893

[11] Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I. S., & Xie, S. (2023). ConvNeXt V2: Co-designing and scaling ConvNets with masked autoencoders. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 16133–16142). IEEE. https://doi.org/10.1109/CVPR52729.2023.01548

[12] Woo, S., Park, J., Lee, J.-Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In V. Ferrari, M. Hebert, C. Sminchisescu, & Y. Weiss (Eds.), Computer vision – ECCV 2018 (Lecture Notes in Computer Science, Vol. 11211, pp. 3–19). Springer. https://doi.org/10.1007/978-3-030-01234-2_1

[13] Su, T., Wu, T., Zhao, J., Scaglione, A., & Xie, L. (2025). A review of safe reinforcement learning methods for modern power systems. Proceedings of the IEEE, 113(3), 213–255. https://doi.org/10.1109/JPROC.2025.3584656

[14] Penso, C., Frenkel, L., & Goldberger, J. (2024). Confidence calibration of a medical imaging classification system that is robust to label noise. IEEE Transactions on Medical Imaging, 43(6), 2050–2060. https://doi.org/10.1109/TMI.2024.3353762

[15] Huang, J., Lou, G., Gu, W., & Shen, C. (2025). Microgrid formation method for load restoration in distribution network with dynamic frequency constraints. Journal of Modern Power Systems and Clean Energy, 13(4), 1323–1334. https://doi.org/10.35833/MPCE.2024.000152

[16] Zhang, Y., Xie, Y., Cai, S., Wu, Q., Zhu, H., & Xiang, Z. (2025). Coordinated control of grid-forming wind turbines and grid-forming energy storage systems for power system restoration. IEEE Transactions on Sustainable Energy, 16(4), 2812–2827. https://doi.org/10.1109/TSTE.2025.3569537

[17] Li, D., Sun, Q., Wang, R., & Sui, Z. (2025). Transient stability analysis and enhancement of inverter-based microgrid considering current limitation. IEEE Transactions on Power Electronics, 40(1), 2429–2441. https://doi.org/10.1109/TPEL.2024.3467215

[18] Karimjee, H., Ranade, S., Ramasubramanian, D., & Lavrova, O. (2025). An analysis of grid-forming inverters in blackstart applications during transformer energization sequence. IEEE Transactions on Energy Conversion, 40(3), 2655–2666. https://doi.org/10.1109/TEC.2025.3567588

[19] Fazlhashemi, S. S., Khodayar, M. E., Khodayar, M., & Wang, J. (2026). Distribution service restoration: A distributionally robust reinforcement learning approach. IEEE Transactions on Smart Grid, 17(1), 690–703. https://doi.org/10.1109/TSG.2025.3597888

[20] Zhang, H., Zhang, Y.-F., Zhang, Z., Wen, Q., & Wang, L. (2024). LogoRA: Local-global representation alignment for robust time series classification. IEEE Transactions on Knowledge and Data Engineering, 36(12), 8718–8729. https://doi.org/10.1109/TKDE.2024.3459908

[21] Pei, J., Wang, J., Shi, D., & Wang, P. (2024). Detection and imputation-based two-stage denoising diffusion power system measurement recovery under cyber-physical uncertainties. IEEE Transactions on Smart Grid, 15(6), 5965–5980. https://doi.org/10.1109/TSG.2024.3407544

[22] Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I. S., & Xie, S. (2023). ConvNeXt V2: Co-designing and scaling ConvNets with masked autoencoders. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 16133–16142). IEEE. https://doi.org/10.1109/CVPR52729.2023.01548

[23] Woo, S., Park, J., Lee, J.-Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In V. Ferrari, M. Hebert, C. Sminchisescu, & Y. Weiss (Eds.), Computer vision – ECCV 2018 (Lecture Notes in Computer Science, Vol. 11211, pp. 3–19). Springer. https://doi.org/10.1007/978-3-030-01234-2_1

[24] Hamel-De le Court, E., Belardinelli, F., & Goodall, A. W. (2025). Probabilistic shielding for safe reinforcement learning. Proceedings of the AAAI Conference on Artificial Intelligence, 39(15), 16091–16099. https://doi.org/10.1609/aaai.v39i15.33767

[25] Shen, F., Zheng, X., Lyu, M., Huang, S., Wu, Q., Shen, Y., Wang, P., & Zhou, Y. (2024). Transactive energy-based sequential load restoration of distribution systems with networked microgrids under uncertainty. IEEE Transactions on Smart Grid, 15(3), 2601–2613. https://doi.org/10.1109/TSG.2023.3322682

[26] Guignard, F., Ginsbourger, D., Levy Häner, L., & Herrera, J. M. (2024). Some combinatorics of data leakage induced by clusters. Stochastic Environmental Research and Risk Assessment, 38(7), 2815–2828. https://doi.org/10.1007/s00477-024-02715-1

[27] Lyu, X., Du, W., Mohiuddin, S. M., Nandanoori, S. P., & Elizondo, M. (2024). Criteria for grid-forming inverters transitioning between current limiting mode and normal operation. IEEE Transactions on Power Systems, 39(4), 6107–6110.

https://doi.org/10.1109/TPWRS.2024.3402012

[28] Wang, P., Li, S., Wang, W., & Xu, D. (2025). A transient stability-enhanced method for grid-forming VSCs considering the current limiting algorithm. IEEE Transactions on Industry Applications, 61(1), 383–394. https://doi.org/10.1109/TIA.2024.3462908

[29] Barbalho, P. I. N., Moraes, A. L., Lacerda, V. A., Barra, P. H. A., Fernandes, R. A. S., & Coury, D. V. (2025). Reinforcement learning solutions for microgrid control and management: A survey. IEEE Access, 13, 39782–39799. https://doi.org/10.1109/ACCESS.2025.3546578

[30] Shao, Z., Wang, F., Xu, Y., Wei, W., Yu, C., Zhang, Z., Yao, D., Sun, T., Jin, G., Cao, X., Cong, G., Jensen, C. S., & Cheng, X. (2025). Exploring progress in multivariate time series forecasting: Comprehensive benchmarking and heterogeneity analysis. IEEE Transactions on Knowledge and Data Engineering, 37(1), 291–305. https://doi.org/10.1109/TKDE.2024.3484454

[31] Zhang, M., Cai, S., Xie, Y., Zhou, B., Zheng, W., Wu, Q., & Wen, J. (2025). Supply resilience constrained scheduling of MEGs for distribution system restoration: A stochastic model and FW-PH algorithm. IEEE Transactions on Smart Grid, 16(1), 194–208. https://doi.org/10.1109/TSG.2024.3446873

[32] Schubert, M. E., Langerman, D., & George, A. D. (2024). Benchmarking inference of transformer-based transcription models with clustering on embedded GPUs. IEEE Access, 12, 123276–123293. https://doi.org/10.1109/ACCESS.2024.3426471

[33] Xiao, Z., Xing, H., Qu, R., Feng, L., Luo, S., Dai, P., Zhao, B., & Dai, Y. (2024). Densely knowledge-aware network for multivariate time series classification. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 54(4), 2192–2204. https://doi.org/10.1109/TSMC.2023.3342640

[34] Hao, G., Li, Y., Li, Y., Jiang, L., & Zeng, Z. (2025). Lyapunov-based safe reinforcement learning for microgrid energy management. IEEE Transactions on Neural Networks and Learning Systems, 36(6), 9985–9999. https://doi.org/10.1109/TNNLS.2024.3496932

[35] Selim, A., Zhao, J., Dong, J., & Lian, J. (2026). Safe deep reinforcement learning for robust frequency and voltage-constrained networked microgrid restoration. IEEE Transactions on Industry Applications, 62(2), 3635–3647. https://doi.org/10.1109/TIA.2025.3626472

[36] Wang, Y., Qiu, D., Teng, F., & Strbac, G. (2024). Towards microgrid resilience enhancement via mobile power sources and repair crews: A multi-agent reinforcement learning approach. IEEE Transactions on Power Systems, 39(1), 1329–1345. https://doi.org/10.1109/TPWRS.2023.3240479

[37] Yang, X., Li, Z., Ding, L., Hu, Z., Fang, H., Sun, W., Wu, Q., & Wen, J. (2025). Reconfigurable converter-formed soft open point-assisted dynamic restoration for distribution systems with cold-load pickup. IEEE Transactions on Sustainable Energy, 16(4), 2271–2285. https://doi.org/10.1109/TSTE.2025.3532911

[38] Pal, A., Hosseinabadi, P. A., Panigrahi, B. K., & Pota, H. R. (2025). A black start solution for voltage-controlled inverters with chattering-free fixed-time sliding mode power synchronization control. IEEE Transactions on Industry Applications, 61(5), 8015–8026. https://doi.org/10.1109/TIA.2025.3559033