

A Human-on-the-Loop Cyber-Physical Framework for Fault-Aware Service Restoration in Nigerian Distribution Feeders: Requirements Derived from EEDC Industrial Training

Okoro Chukwueloka Nnaemeka

Abstract—Distribution fault management depends on more than fault recognition. An operator or automation controller must establish the affected section, confirm switching-device states, find a feasible post-fault topology, respect protection and equipment limits, and verify that service has returned without creating a new hazard. This paper proposes a human-on-the-loop cyber-physical architecture for distribution feeders in Nigeria, with its requirements grounded in the author's industrial training at the Enugu Electricity Distribution Company (EEDC)/MainPower. The training report records field exposure to right-of-way hazards, sagging conductors, transformer load assessment and commissioning, fuse-rating checks, relay work, insulation-resistance tests, earthing tests, power-quality inspection, and materials control. These activities motivate a design in which local protection remains independent of machine-learning inference and remote communications, while software proposes only switching actions that pass explicit electrical, topology, state-freshness, and authorization checks. The paper reconciles prior drafts against verified literature and separates published findings from author-proposed methods and claims that lack reproducible evidence. It formulates a constrained feeder-restoration problem, describes a candidate waveform classifier and policy-learning layer, and defines reliability, power-quality, latency, safety, cybersecurity, and economic evaluation measures. The proposed research has not yet been validated on an EEDC feeder; therefore, no numerical improvement in restoration time, SAIDI, SAIFI, or line losses is claimed. A reproducible evaluation plan using OpenDSS and a documented benchmark network is specified before any utility-specific calibration or field pilot. The result is a research framework and validation protocol, not a report of completed field or simulation experiments.

Index Terms—Distribution automation, fault detection isolation and service restoration, human-on-the-loop control, Nigerian electricity distribution, operational technology security, feeder reconfiguration, reliability indices, cyber-physical systems.

I. INTRODUCTION

An electricity distribution feeder has to deliver power within thermal, voltage, and protection limits while accommodating faults, maintenance work, and changing demand. A permanent fault on a radial section may disconnect customers beyond the fault until a crew identifies the affected span, isolates it,

and restores healthy sections where an alternate path exists. The sequence involves measurements, protection, switch status, network topology, work procedures, and operational authority. Fault detection alone does not establish that a proposed switching action is safe. Reviews of fault detection, isolation, and service restoration (FDIR) describe this as a coordinated problem involving detection, faulted-section isolation, and supply restoration [1]–[3].

The motivation for this work developed during the author's industrial training at EEDC/MainPower. The SIWES report records rotations through health, safety and environment (HSE); Stores; Network Expansion; Protection Control and Metering; Electric Filters/Power Quality; Cables; and Maintenance. The author took part in field inspections that recorded bushy substations, right-of-way encroachment, sagging conductors, and damaged equipment in areas including Abakpa and Emene. Other recorded activities include supervised transformer load assessment, observation of transformer installation and commissioning, fuse-rating checks, relay calibration support, insulation-resistance and earth-resistance tests, feeder or line testing, and power-quality inspection [4]. These are observations from an industrial attachment, not a statistically representative survey of Nigerian utilities. They are used here to define practical requirements, not to infer national failure rates.

The report also records materials receiving, stock-card updates, and issue documentation in Stores. This administrative work has a direct engineering implication: a fault can be located quickly while replacement materials or a crew remain unavailable. Likewise, the HSE rotation shows why automation cannot issue commands without knowing whether work is in progress, a permit is active, or a switching boundary has been secured. Network expansion activities exposed the author to load assessment, transformer siting, and fuse selection. Protection Control and Metering activities brought attention to protection settings, current-transformer tests, insulation condition, and earthing. Power-quality and cable work add requirements for voltage and current measurement, fault evidence, and asset condition. The architecture in this paper follows those operational dependencies.

Nigerian distribution research supports a focus on feeder topology and outage diagnosis. Onodjohwo *et al.* review network reconfiguration as a method for addressing losses and voltage performance in Nigerian distribution systems [5]. Akanni *et al.* analyze outage data obtained from Eko Electricity

The author is with the Department of Electronic Engineering, Federal University of Technology, Owerri, Nigeria (e-mail: eloka.okoro.researcher@outlook.com). The field observations summarized here are drawn from the author's SIWES technical report on the Enugu Electricity Distribution Company (EEDC)/MainPower and Federal Radio Corporation of Nigeria (FRCN)/Tait Radio Academy attachment, November 2025. This is a design and methods paper; no utility operational data or executable simulation model was supplied for this study.

Distribution Company and Ikeja Electric and report transient faults and jumper cuts among the leading causes within their studied records [6]. These findings establish a documented local research basis for studying outage response; they do not supply EEDC/MainPower feeder parameters or authorize transfer of results between utilities.

This paper asks four research questions: (1) which automation requirements follow from the author's documented field activities and from the FDIR literature? (2) how should a controller restrict switching actions to electrically feasible, topologically valid, and authorized operations? (3) which measurements would demonstrate a benefit without conflating classifier speed with end-to-end restoration time? and (4) how can the proposal be tested reproducibly before a field trial? The paper makes four contributions. It reconciles earlier manuscript claims against evidence supplied or verifiable in the literature; maps industrial-training observations into requirements and acceptance tests; defines a constrained, human-supervised restoration architecture and mathematical formulation; and supplies a benchmark-based evaluation protocol. The paper does not claim that a trained model, OpenDSS simulation, hardware prototype, or field pilot has already been completed.

II. EVIDENCE BASIS AND RELATED RESEARCH

A. Source selection and claim categories

The evidence set was assembled from the materials supplied for this task: successive versions of the HOTL paper, three Elicit spreadsheets, the SIWES report, the embedded-systems design notes, and the writing-rules document. The citations were checked against DOI records, publisher pages, standards-body pages, or the original institutional source where accessible. The Elicit lists contain useful references but also records with incomplete or generic metadata, including an item titled "Untitled" and an industry paper with an author field rendered as "19". Such entries are not used as factual support. Claims in the former drafts were grouped into four classes: (a) supported by a traceable publication or official document; (b) documented as personal training experience; (c) proposed here as an engineering design; or (d) unresolved because the supplied materials do not contain data or a reproducible derivation.

This is a focused engineering synthesis, not a systematic review. The search was not preregistered, and the included papers were not scored using a formal quality instrument. The source-selection procedure therefore supports design decisions and a validation agenda, but it does not support a quantitative meta-analysis or a claim that all relevant studies have been captured.

B. Distribution faults and service restoration

Srivastava *et al.* organize the FDIR problem and discuss distribution automation, detection methods, and restoration constraints [1]. Arefifar *et al.* review self-healing in modern distribution systems and explain the dependence on sensing, communications, automation, and network configuration [2]. Shen *et al.* group service-restoration approaches into centralized, distributed, and hierarchical arrangements, with methods that include expert rules, graph algorithms, mathematical

programming, and multi-agent systems [3]. This classification informs the proposed hierarchical design: a local relay clears a fault, an edge controller confirms and isolates the section when device logic and interlocks permit, and a supervisory optimizer ranks restoration actions that require broader feeder state.

Research specific to Nigeria provides context but must be used at the right scale. The review by Onodjohwo *et al.* addresses reconfiguration in Nigerian distribution systems, including the effects of radial feeders and the use of optimization methods [5]. Akanni *et al.* base their Lagos outage analysis on fault records from two distribution companies rather than on a synthetic simulation alone [6]. Nwajuonye *et al.* discuss distributed generation as a route to reliability and efficiency improvement in Nigerian networks [7]. These studies motivate a combination of topology, outage evidence, and distributed-resource constraints; none gives permission to describe a proposed EEDC/MainPower model as a measured field result.

C. Reconfiguration and reliability measures

Baran and Wu's distribution-network reconfiguration work formalizes the use of switch exchanges to reduce losses and improve load balancing while retaining a radial configuration [8]. The mathematical model below follows that general problem structure but extends the objective to include unsupplied energy, voltage deviation, switching wear, and a penalty for unsafe or unsupported actions. Reliability reporting uses the definitions in IEEE 1366-2022 rather than values reproduced from earlier drafts [9]. This matters because SAIDI and SAIFI depend on the customer population, interruption records, reporting interval, and inclusion rules. A number without those definitions is not suitable for comparing utilities or claiming improvement.

D. Learning, digital twins, and control authority

Al-Saadi *et al.* survey reinforcement-learning approaches for power management in distribution applications and identify open issues associated with control objectives, training environments, and practical deployment [10]. Schulman *et al.* introduce proximal policy optimization (PPO), a policy-gradient method with a clipped surrogate objective [11]. The present paper treats PPO or multi-agent PPO as optional candidate policies for ranking candidate switching plans, not as a safety layer. A policy trained on simulation may have unrecognized failure cases; deterministic constraints and independent electrical protection must govern whether an action can be executed.

Federated learning can aggregate model updates without centralizing raw samples. McMahan *et al.* introduced a communication-efficient method now widely known as federated averaging [12]. This property may help where utility datasets cannot be moved freely, but it does not eliminate dataset bias, secure aggregation requirements, or the need to validate a model at each deployment location. Sifat *et al.* review electric-grid digital-twin concepts for state representation and analysis [13]. In this design, the digital twin is an engineering model that receives telemetered switch status, feeder measurements, and validated asset parameters. A stale or inconsistent twin is treated as a reason to block automatic switching, rather than as a substitute for field confirmation.

E. Operational technology safety and security

The Nigerian Electricity Regulatory Commission publishes the Distribution Code and a Health and Safety Code for the Nigerian Electricity Supply Industry [14], [15]. These documents are regulatory context for the operating utility; this paper does not interpret them as granting authority for autonomous switching. IEEE 1366-2022 provides a recognized basis for distribution reliability indices [9]. NIST SP 800-82 Rev. 3 describes security considerations for operational technology (OT), including systems that monitor or directly affect physical processes [16]. These sources support local protection independence, role-based access, network segmentation, event logging, and a controlled response to communications or data-integrity failures.

III. FIELD-DERIVED REQUIREMENTS FROM EEDC/MAINPOWER

The author's SIWES report covers work performed under supervision and does not imply authority to operate live equipment independently. The relationship between the experience and the proposed system is stated in Table II. The table is a traceability device: a requirement is included only where a recorded activity or a well-established FDIR need gives it a concrete engineering reason.

Three aspects of the training shaped the design. First, protection operations need conservative authority boundaries. A software classifier can assist diagnosis, but it must not replace a properly configured relay or inhibit a trip. Second, the evidence required to justify a switching action extends beyond a predicted fault label. Position indication, feeder loading, phase voltages, equipment ratings, interlocks, and work status must all be known or conservatively bounded. Third, utility work includes physical hazards and logistics. The system must record who owns the next action, what has been verified, and whether a field crew remains necessary.

The paper therefore uses the phrase *human-on-the-loop* in a narrow operational sense: an authorized operator can observe system state, inhibit automation, change the operating mode, and require manual authorization for actions outside a predefined envelope. The concept is not permission for the software to bypass standing operating procedures. Immediate fault clearing remains local to existing protection. Any implementation must be approved by the utility and adapted to its switching rules, protection coordination, cybersecurity controls, and permit-to-work process.

IV. PROPOSED CYBER-PHYSICAL ARCHITECTURE

The architecture has four functional layers: physical measurement and switching; local protection and edge processing; supervisory communication and state estimation; and operator-facing decision support. These are logical layers, not a procurement bill of materials. A small pilot may use existing relays, recloser controllers, meters, and SCADA points rather than introduce a new sensor at every location.

A. Measurement and event acquisition

Candidate inputs include phase voltages and currents, residual or zero-sequence quantities where available, frequency, breaker/recloser position, relay targets, trip timestamps, feeder loading, transformer ratings, and communication-health flags. Sensor sampling must match the use case. Protection-quality waveforms generally require local sampling and processing that is separate from slower supervisory telemetry. The paper does not prescribe sample rates or sensor classes without the existing instrument-transformer and relay specifications.

Each event record should include an asset identifier, the time source and synchronization quality, measurement-quality flags, the current switch-state snapshot, relay targets, model version, and operator or automation actions. A missing value must not be treated as a measured zero. Stale states should carry age metadata. The state estimator may fuse SCADA measurements with topology and load estimates, but the switching controller must preserve the difference between measured, estimated, and unknown states.

B. Local processing and communication

The first release should use protection and deterministic interlocks that already have an approved role in the scheme. An edge computer may form event windows, reject corrupted or saturated signals, run a candidate classifier, and transmit a fault hypothesis. It should not issue a breaker-close command solely because a neural network labels an event as a transient fault. For permanent faults, relay protection or an approved local scheme clears the fault first; restoration logic then checks isolation and alternate-feed constraints.

The communication design must tolerate interruptions and delay. Fibre, utility radio, licensed wireless, cellular, or low-bandwidth telemetry can be assessed against coverage, latency, service availability, operating cost, and cybersecurity needs. This paper does not assume that one communications technology is available throughout an EEDC feeder. Telemetry queues should retain sequence numbers and timestamps, reject replayed commands, and mark old messages as stale. A command that cannot be acknowledged leaves the controller in a pending or blocked state; the controller does not repeat a close command blindly.

C. Digital twin and operator service

The digital twin stores the feeder graph, line impedances, transformer ratings, loads, phase connections, switch states, protection zones, and approved operating limits. Its main purpose in the first release is to test a candidate action against the last verified network state. Prediction of future asset condition or thermal capacity can be added when there are measured asset parameters and a method for uncertainty bounds. A twin is not a second physical protection system.

The operator display should answer five questions without requiring the operator to infer the model's intent: What alarm occurred? Which measurements support the fault hypothesis? Which devices are confirmed open or closed? Which constraints were checked? Why was an action allowed or rejected? Each proposed sequence includes an expected post-switching state,

TABLE I
SELECTED LITERATURE AND ITS USE IN THE PROPOSED STUDY

Source	Evidence type	Topic established by source	Use and boundary in this paper
Akanni <i>et al.</i> [6]	Field-record analysis	Outage causes in Lagos using fault reports from EKEDC and IKEDC	Local context for fault reporting; not treated as EEDC/MainPower data.
Onodjohwo <i>et al.</i> [5]	Review	Nigerian distribution reconfiguration and voltage/loss concerns	Supports network reconfiguration as a research direction; no result imported as this paper's performance.
Srivastava <i>et al.</i> [1]	Review	FDIR methods and distribution automation	Structures detection, isolation, and restoration stages.
Arefifar <i>et al.</i> [2]	Review	Self-healing system components and open requirements	Supports layered sensing and automation with network constraints.
Shen <i>et al.</i> [3]	Review	Centralized, distributed, and hierarchical restoration methods	Motivates local-plus-supervisory control.
Baran and Wu [8]	Optimization method	Feeder switching for loss reduction and load balancing	Starting point for topology and loss equations; updated objective is proposed here.
Al-Saadi <i>et al.</i> [10]	Survey	Reinforcement learning for distribution power management	Positions RL as a research candidate requiring safety constraints.
McMahan <i>et al.</i> [12]	Learning method	Federated averaging for distributed data	Supports a possible future data-governance route, not a required first release.
Sifat <i>et al.</i> [13]	Review	Digital-twin concepts for power grids	Informs the model-to-measurement synchronization layer.
NIST SP 800-82 Rev. 3 [16]	Security guidance	OT architecture, threats, and safeguards	Informs segmentation, access control, event logging, and degraded modes.




Fig. 1. Proposed layered architecture. Vertical arrows show the primary decision pipeline. Status and feedback are exchanged between adjacent service layers; local relay protection remains independent of the supervisory service.

affected customer or load records if available, confidence and data-age indicators, and a rollback or safe-state action. A decision that cannot be explained from the recorded inputs should not be eligible for automatic execution.

D. Control modes and authority

Three operating modes are proposed. In *advisory mode*, the system analyzes events and displays plans but cannot command switches. In *supervised mode*, the system may execute only pre-approved actions inside a narrow operating envelope while the dispatcher can block the sequence. In *manual/fallback mode*, automatic switching is inhibited and the operator follows existing procedures. Mode changes are authenticated and logged. Loss of communication, invalid topology, stale switch state, detected cyber-integrity failure, or a failed interlock moves the control function to the blocked or fallback state, not to a more autonomous mode.

V. MATHEMATICAL FORMULATION

A. Network model

Represent a feeder by a graph $G = (\mathcal{N}, \mathcal{E})$ with source nodes $\mathcal{S}$, loads $\mathcal{L}$, controllable switches $\mathcal{K}$, and candidate tie branches. The binary variable $u_{ij} \in \{0, 1\}$ indicates whether branch (i, j) is closed. The vector $\mathbf{u}^0$ records the last verified configuration; $\mathbf{u}$ is a proposed configuration. A node status $z_i \in \{0, 1\}$ indicates whether node i is energized. A restoration decision cannot energize an unserved node unless its path reaches an authorized source or an explicitly approved islanded source.

Let P_{ij} and Q_{ij} denote branch active and reactive power flow; $v_i = |V_i|^2$ is squared voltage magnitude; and $\ell_{ij} = |I_{ij}|^2$ is squared branch current magnitude. For a radial branch-flow approximation, the steady-state constraints may be written as

$$P_{ij} = \sum_{k:(j,k) \in \mathcal{E}} P_{jk} + p_j + r_{ij}\ell_{ij}, \quad (1)$$

$$Q_{ij} = \sum_{k:(j,k) \in \mathcal{E}} Q_{jk} + q_j + x_{ij}\ell_{ij}, \quad (2)$$

TABLE II
TRACE FROM DOCUMENTED INDUSTRIAL TRAINING TO DESIGN REQUIREMENTS

Training area	Activity recorded in SIWES report	Requirement for the proposed system	Acceptance evidence
HSE	Inspections of bushy substations, right-of-way encroachment, sagging wires, and damaged equipment	Report hazards as asset/work records; prohibit remote control when work permits, earthing, or safe access status are unresolved	Every hazardous or unknown work-state blocks an affected switching plan; the event is logged for a human decision.
Stores	Count and inspect received materials; update stock-control cards and issue documents	Link the control-room event to a work order and parts status; do not equate electrical restoration with completion of physical repair	Work-order tests show correct stock/status trace and an explicit “crew/material pending” state.
Network Expansion	Load assessment, transformer siting, installation observation, completion checklist, and fuse-rating calculations	Store feeder connectivity, transformer ratings, load envelopes, and approved protection settings as versioned model data	A load-flow case reproduces the approved baseline within a stated engineering tolerance; changes are auditable.
Protection Control and Metering	Relay calibration support, transformer insulation-resistance testing, earth-resistance tests, and line testing	Respect relay trips and lockouts; check freshness and plausibility of current, voltage, breaker, and earthing/work-state information	Hardware-in-the-loop or replay tests prove the optimizer cannot bypass relay trips, lockouts, or an active work block.
Electric Filters / Power Quality	Incomer-panel checks and attention to harmonics and voltage fluctuations	Use voltage magnitude, phase imbalance, harmonics where instrumentation supports them, and quality flags as constraints or alarm variables	Metric reports retain phase, sampling interval, sensor quality, and limits; missing phase data blocks unsupported switching.
Cables	Cable types, installation practices, and fault-finding procedures	Preserve asset segment identity and fault location uncertainty; avoid declaring a physical cable repaired from remote waveform evidence alone	Simulated and replayed events report the estimated section and uncertainty; repair confirmation remains a field task.
Maintenance	Preventive and corrective routines on transformers, switchgear, and associated assets	Track device health, operation counts, maintenance lockouts, and failed-switch behavior	Tests cover an unavailable switch, stale position indication, operation-limit exceedance, and maintenance mode.
FRCN/Tait Academy	Radio Exposure to radio systems, transmitters, receivers, antenna systems, and communication protocols	Treat the communications path as potentially delayed or unavailable; retain a local protective mode and bounded data buffering	Link interruption, packet delay, and recovery tests do not affect local protection or create an unverified close command.

$$v_j = v_i - 2(r_{ij}P_{ij} + x_{ij}Q_{ij}) + (r_{ij}^2 + x_{ij}^2)\ell_{ij}, \quad (3) \quad \text{where}$$

$$P_{ij}^2 + Q_{ij}^2 \leq v_i \ell_{ij}. \quad (4)$$

Here p_j and q_j are net downstream demand under the selected sign convention, while r_{ij} and x_{ij} are branch resistance and reactance. Equation (4) is a second-order-cone relaxation of the branch current relation. A candidate obtained from a relaxation requires validation using a full AC load flow before a switching command is accepted. For three-phase feeders, these equations must be replaced or checked with an unbalanced multiphase model. A balanced single-phase equivalent is inadequate where phase imbalance controls the operating limit.

B. Objective function

The action-selection problem minimizes a weighted cost over a proposed topology, load allocation, and service status:

$$\min_{\mathbf{u}, \mathbf{z}, \mathbf{p}} J = w_L P_{\text{loss}} + w_V D_V + w_E E_{\text{NS}} + w_S N_{\text{sw}} + w_R R_{\text{op}} + w_U U_{\text{unc}}, \quad (5)$$

$$P_{\text{loss}} = \sum_{(i,j) \in \mathcal{E}} r_{ij} \ell_{ij}, \quad (6)$$

$$D_V = \sum_{i \in \mathcal{N}} z_i \left(\frac{|V_i| - V_i^{\text{ref}}}{V_i^{\text{ref}}} \right)^2, \quad (7)$$

$$E_{\text{NS}} = \sum_{i \in \mathcal{L}} P_i^{\text{crit}} (1 - z_i) \Delta t, \quad (8)$$

$$N_{\text{sw}} = \sum_{(i,j) \in \mathcal{K}} |u_{ij} - u_{ij}^0|. \quad (9)$$

P_{loss} is active technical loss; D_V penalizes voltage deviation; E_{NS} is an interval estimate of unsupplied energy; N_{sw} measures switching actions; R_{op} is a risk score tied to approved risk controls; and U_{unc} penalizes data uncertainty or model mismatch. The weights w must be selected with the utility and reported with units or a normalization rule. They cannot be tuned after observing test outcomes without disclosing the procedure. Safety constraints are hard constraints and cannot be traded for a lower weighted objective.

C. Electrical and topology constraints

The planner must maintain voltage and current within approved operating ranges:

$$V_i^{\min} \leq |V_i| \leq V_i^{\max}, \quad i \in \mathcal{N} : z_i = 1, \quad (10)$$

$$|I_{ij}| \leq I_{ij}^{\max}, \quad (i, j) \in \mathcal{E} : u_{ij} = 1, \quad (11)$$

$$|S_{ij}| \leq S_{ij}^{\max}, \quad (i, j) \in \mathcal{E} : u_{ij} = 1. \quad (12)$$

Switching must not create an unauthorized cycle. For a single-source radial network with no intentional islands, the energized closed-branch graph must be connected to its source and acyclic, which implies $|\mathcal{E}_{\text{closed}}| = |\mathcal{N}_{\text{energized}}| - 1$. The edge-count condition alone is insufficient: connectivity and acyclicity must both be checked. A network with multiple permitted sources must instead form an authorized rooted forest, with each energized component containing exactly one designated source unless a documented protection scheme permits a different structure. A graph traversal or mixed-integer flow constraints can test these properties before the AC load-flow check.

Three-phase imbalance can be represented by the voltage-unbalance factor

$$\text{VUF}_i = 100 \frac{|V_i^-|}{|V_i^+|}, \quad (13)$$

where V_i^- and V_i^+ are negative- and positive-sequence voltage components. The admissible threshold is a utility-approved parameter; this paper does not impose an arbitrary universal value. Protection coordination also depends on actual relay curves, fault current, fuse ratings, and directional or distributed-generation effects. The planner must call a protection-coordination check using the approved data. It must not adopt a fixed coordination time from an earlier draft without a settings study.

D. Human authorization and data validity constraints

Define a binary data-validity flag d_k for each critical observation, an age a_k for its timestamp, a confidence score c_f for an optional classifier, and a work-block flag b_i for each affected asset or section. A candidate close action at switch j is eligible only if

$$\begin{aligned} \text{Eligible}(j) = & d_j \mathbb{I}[a_j \leq a_{\max}] \mathbb{I}[c_f \geq \theta_f] \\ & \times (1 - b_j) \mathbb{I}[\text{electrical} = 1] \\ & \times \mathbb{I}[\text{topology} = 1] \\ & \times \mathbb{I}[\text{authority} = 1] = 1. \end{aligned} \quad (14)$$

The indicator for classifier confidence is omitted where the classifier is not used or where the protection scheme does not require it. It never replaces relay logic. A failed or unknown factor sets eligibility to zero. Thresholds $a_{\max}$ and θ_f are to be established through the validation process and device requirements, not guessed in the paper.

E. Candidate waveform classifier

A neural model may rank fault hypotheses from a local waveform window $X \in \mathbb{R}^{T \times F}$ containing T samples and F

channels. A candidate architecture combines temporal convolutions and bidirectional long short-term memory (BiLSTM) features with attention. For convolutional stage m ,

$$Z^{(m)} = \phi \left(W^{(m)} * Z^{(m-1)} + b^{(m)} \right), \quad Z^{(0)} = X, \quad (15)$$

where $*$ denotes a one-dimensional convolution and ϕ is an activation function. The bidirectional recurrent state at time t is

$$\begin{aligned} H_t &= [\vec{h}_t; \overleftarrow{h}_t], \\ \vec{h}_t &= \text{LSTM}_f(Z_t), \quad \overleftarrow{h}_t = \text{LSTM}_b(Z_t). \end{aligned} \quad (16)$$

Attention weights and a context vector can be defined as

$$e_t = q^\top \tanh(W_h H_t + b_h), \quad (17)$$

$$\alpha_t = \frac{\exp(e_t)}{\sum_{\tau=1}^T \exp(e_\tau)}, \quad c = \sum_{t=1}^T \alpha_t H_t, \quad (18)$$

$$\hat{y} = \text{softmax}(W_o c + b_o). \quad (19)$$

This architecture is a candidate, not a selected or trained model. A simpler feature-based classifier may be preferable when labeled events are few or instrumentation differs by feeder. Evaluation must use time- and feeder-separated train, validation, and test sets to reduce leakage between similar waveform windows. Relevant outputs include per-class precision and recall, confusion matrix, false-trip proxy rate, calibration error, out-of-distribution performance, memory footprint, and 95th-percentile inference latency. Synthetic waveforms must be distinguished from field-labelled examples. A high aggregate accuracy does not establish readiness for protection use.

F. Policy learning as a non-safety-critical planner

If a reinforcement-learning agent is evaluated, its state s_t may include topology, loading, voltage margins, device availability, fault location uncertainty, communications status, and remaining operation budgets. Its action a_t is a proposal from a finite candidate set of switch sequences. A reward can be based on the negative constrained cost in (5), with large penalties for violations during training. However, reward penalties are not a safety proof. The deployable action is the projection of a candidate onto the certified feasible set:

$$a_t^{\text{exec}} = \begin{cases} a_t^{\text{proposal}}, & a_t^{\text{proposal}} \in \mathcal{A}_{\text{safe}}(s_t), \\ \emptyset, & \text{otherwise,} \end{cases} \quad (20)$$

where $\mathcal{A}_{\text{safe}}(s_t)$ contains only actions that pass state, topology, electrical, protection, permit, and authorization checks. The empty action means no switching command. PPO's clipped objective is

$$L^{\text{CLIP}}(\theta) = \mathbb{E}_t \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right], \quad (21)$$

where $r_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{\text{old}}}(a_t|s_t)$ and $\hat{A}_t$ is an advantage estimate [11]. An RL model is not necessary for the minimum viable system. Rule-based constraints and mixed-integer or heuristic optimization should be established as baselines first.

Algorithm 1 Constrained restoration supervisor

Require: Event e , verified topology G , approved settings Π , work blocks B

- 1: Receive relay event; preserve raw measurements and timestamps
- 2: **if** relay state, switch position, time quality, or topology is invalid **then**
- 3: Block automatic switching; alarm and request confirmation
- 4: **return**
- 5: **end if**
- 6: Estimate faulted section and uncertainty set $\mathcal{F}$
- 7: Confirm isolation using relay targets and available position feedback
- 8: **if** fault isolation is not confirmed **then**
- 9: Keep restoration blocked; request dispatcher action
- 10: **return**
- 11: **end if**
- 12: Generate candidate configurations $\mathcal{C}$ using approved switches
- 13: **for all** $c \in \mathcal{C}$ **do**
- 14: Check connectivity, radiality, ratings, voltage, phase imbalance, and protection
- 15: Reject c if any hard constraint or work block fails
- 16: **end for**
- 17: **if** $\mathcal{C}$ is empty **then**
- 18: Record rejection reasons and present manual options
- 19: **return**
- 20: **end if**
- 21: Rank feasible candidates by objective and predicted customer interruption
- 22: **if** mode permits execution and selected plan is pre-authorized **then**
- 23: Execute one operation at a time; await positive acknowledgement
- 24: Re-estimate state and re-check constraints after each operation
- 25: **else**
- 26: Present plan, evidence, and constraints for dispatcher decision
- 27: **end if**
- 28: Verify post-action electrical state; log outcome and model versions

VI. EVENT HANDLING AND SWITCHING LOGIC

Fig. 2 gives the event sequence. The sequence separates immediate protection from post-fault service restoration. It also makes a failed precondition a first-class result, not an exception that triggers a risky best-effort close.

A. Isolation and restoration are different tasks

A relay trip is evidence that the protective scheme operated; it is not proof that the failed segment is fully isolated. A valid state may require current-transformer measurements, relay targets, two-end measurements, switch auxiliary contacts, or a field confirmation, depending on the feeder. If a switch

position is unknown, the controller must not assume the commanded state became the physical state. It should wait for a positive acknowledgement within the approved timeout, use independent measurements where available, and block downstream switching if the state remains uncertain.

The planner evaluates a candidate sequence as a complete transition, including the state after each operation. For example, a tie switch should not close into a faulted section; the relevant sectionalizing boundaries must be open and positively confirmed first. A sequence is rejected if closing the tie overloads a transformer or line, violates voltage limits, creates an unapproved loop, exposes a maintenance worksite, or invalidates relay coordination. Once a switch operates, the state estimate must be refreshed before a second dependent command is issued.

B. Safe-state and rollback semantics

Rollback does not mean blindly returning to the prior switching vector. If a new switch operation fails or a voltage constraint is violated, the controller must execute only a previously reviewed and feasible recovery sequence. Where the current physical state is uncertain, the controller stops issuing new commands and transfers control to the dispatcher. This rule prevents an automated rollback from recreating a fault path or energizing equipment under work. An emergency trip by the protective relay remains independent of supervisory rollback logic.

C. Operator workload and decision records

The system should present ranked actions and their reasons rather than a stream of unprioritized telemetry. The event record contains the event identifier, asset identifiers, measurement quality, current and proposed switch vectors, power-flow result, constraint checks, protection-study version, communication status, operator identity, approvals, command identifiers, acknowledgements, and final measurements. The operator can compare feasible options against customer criticality and crew constraints where reliable records exist. Any weighting of critical load categories requires the utility to define and approve the categories. An inferred customer class must not silently replace an approved priority register.

VII. REPRODUCIBLE EVALUATION PROTOCOL

The proposed architecture has not been run against an EEDC/MainPower feeder in the supplied material. Evaluation must therefore be staged. The first stage establishes the method on a recognized radial benchmark with published parameters. The IEEE 33-bus distribution test system is a suitable starting point for method development; its published parameters should be transcribed into the simulator and independently checked before experiments. A utility-specific test should follow only after the utility approves data access and validates topology, impedances, load estimates, phase connections, ratings, and protection settings. The EPRI OpenDSS distribution simulator provides a documented tool for time-series and distribution-system studies [17].




Fig. 2. Proposed FDIR decision sequence. The left branch blocks switching when the state cannot be verified. The right branch checks isolation and all hard constraints before execution; failed checks return the process to a blocked state.

A. Model and data preparation

The model package should identify simulator and software versions, circuit data source, load model, per-phase connections, line impedance units, transformer configuration, regulator and capacitor controls, source voltage, and solver tolerances. Each modelling assumption should be recorded in a parameter register. At least one independent engineer should review node naming, branch connectivity, transformer ratios, line units, and base values before simulation. The validated base case must reproduce a power-flow solution under stated tolerances before any fault or reconfiguration experiment begins.

Synthetic scenarios should include single-line-to-ground, line-to-line, double-line-to-ground, three-phase, transient, permanent, high-impedance where the model supports it, switch failure, bad position feedback, missing measurements, delayed communications, phase imbalance, overloaded tie paths, and active maintenance blocks. Fault impedance and location should be varied over a documented range; random seeds should be saved. If waveform data are generated from a simulator, it should be labelled synthetic and not merged invisibly with field records. A field dataset requires permission, de-identification as appropriate, event labels reviewed by engineers, time-quality assessment, and an explicit policy for classifying ambiguous events.

B. Baselines and ablations

At least four baselines are required: (B1) local protection plus manual restoration; (B2) a deterministic rule-based FDIR sequence; (B3) constrained mathematical or heuristic reconfiguration without learned classification; and (B4) the complete candidate system with the same hard safety filter. An optional fifth baseline can include the learned classifier but no learned control policy. Ablations should separately remove the classifier, freshness check, topology check, digital-twin power-flow check, and operator authorization gate. The purpose is to identify which component changes diagnostic quality, restoration time, or safety behavior. Every baseline must use the same disturbance set, initial state, protection settings, load data, and scoring rules.

C. Metrics and definitions

For a customer population N_T , interruption events i , customers interrupted N_i , and interruption duration U_i , the standard reliability measures include

$$\text{SAIDI} = \frac{\sum_i N_i U_i}{N_T}, \quad (22)$$

$$\text{SAIFI} = \frac{\sum_i N_i}{N_T}, \quad (23)$$

$$\text{CAIDI} = \frac{\text{SAIDI}}{\text{SAIFI}}, \quad \text{when SAIFI} > 0. \quad (24)$$

These equations provide a simplified event-based description; the applicable IEEE 1366 definitions and reporting rules govern

TABLE III
EVALUATION MATRIX AND REQUIRED EVIDENCE

Test family	Fault or condition	Primary measures	Evidence required before reporting a result
Classifier quality	Fault type, location uncertainty, noisy or missing channels	Per-class precision/recall, confusion matrix, calibration, false-positive rate	Label source, train/test separation, seeds, waveform provenance, class counts, confidence intervals.
Restoration	Isolatable permanent fault with healthy alternate feed	End-to-end restoration time, ENS, loads restored, operations count	Start/stop event definitions, switch sequence and acknowledgements, current/voltage after each action.
Electrical feasibility	High loading, voltage sag, phase imbalance, reverse power where applicable	Voltage extrema, current/thermal violations, losses, protection-coordination pass rate	AC power-flow output, line and transformer limits, per-phase model, approved thresholds.
Topology integrity	Tie closure, switch failure, incorrect position indication	Cycle detection, energized-source connectivity, blocked unsafe plan rate	Graph state before/after, independent radiality/connectivity tests, switch-position evidence.
Communication degradation	Link loss, delay, stale packets, replay attempt	Blocked-command rate, stale-state detection, recovery behavior	Packet trace, injected fault configuration, event timestamps, authentication log.
Worksite and HSE	Active permit, maintenance lockout, unknown earth/isolator state	Unauthorized-action count, correct inhibition, alert latency	Permit/work-block record, test scenario and command audit.
Reliability and losses	Repeated interruption and load variation over the same interval	SAIDI, SAIFI, CAIDI, ENS, real losses, voltage profile	Customer counts, interruption inclusion rules, simulation duration, base case, confidence or dispersion.
Runtime	Increasing feeder size and candidate count	Median and 95th-percentile decision latency, memory, timeout frequency	Hardware specification, software versions, run count, warm-up and timing boundaries.

the final calculation, including treatment of planned events and major event days. A simulation should report its simulated population and interval, not label its numbers as annual utility indices unless there is a defensible temporal and statistical mapping. The percentage reduction in a metric is

$$\Delta\% = 100 \frac{M_{\text{base}} - M_{\text{method}}}{M_{\text{base}}}, \quad (25)$$

where M_{base} and M_{method} come from matched experiments. Both values and the experimental sample size should accompany the percentage.

Restoration time must be decomposed into detection, confirmation, isolation, planning, command, and verification stages:

$$T_{\text{rest}} = T_{\text{det}} + T_{\text{conf}} + T_{\text{iso}} + T_{\text{plan}} + T_{\text{cmd}} + T_{\text{ver}}. \quad (26)$$

A classifier's inference latency is only T_{det} or part of it; it is not total restoration time. Timing should use synchronized monotonic clocks where possible. Report median, 95th percentile, maximum, and the number of cases in which no safe automatic restoration was found. A safe abstention is a valid outcome and should not be scored as a successful restoration.

For classifier assessment, report precision $TP/(TP + FP)$, recall $TP/(TP + FN)$, and per-class confusion counts. For safety assessment, count prohibited commands proposed, prohibited commands accepted by the safety filter, and unauthorized physical operations. The required number of accepted prohibited commands is zero in the tested scenarios; a zero count in a finite test set is not proof that the probability is universally zero. State uncertainty and false negatives should be reported alongside a fault classifier's speed. Power-quality outcomes include minimum and maximum phase voltage, voltage-unbalance factor, line and transformer loading, and any approved harmonic metric supported by the available instrumentation.

D. Statistical comparison

The experimental design should use paired scenarios. For a given scenario s , the difference between the baseline and proposed method is $d_s = M_{s,\text{base}} - M_{s,\text{proposed}}$. Report the mean and median of d_s , its dispersion, and a bootstrap confidence interval or a justified parametric interval. Report failure cases individually when the sample is small. Repeated runs with the same random seed do not count as independent disturbance cases. If learned models use random initialization or stochastic policies, disclose the number of independent training runs and the selection rule used to choose a checkpoint.

E. Reproducibility package

The planned package should contain (1) circuit model and data provenance; (2) input and parameter register; (3) versioned scripts or notebooks; (4) dependency and simulator versions; (5) fault scenario generator and random seeds; (6) trained model weights and preprocessing steps if learning is used; (7) baseline definitions; (8) raw outputs and summary calculations; (9) event logs and rejection reasons; and (10) a README with one command or one documented sequence to reproduce each table. Utility data remain subject to the utility's security and privacy rules. If those data cannot be shared, provide an anonymized synthetic case and describe exactly which results cannot be independently verified.

VIII. RECONCILIATION OF EARLIER DRAFTS

The earlier versions contain a consistent research direction but mix proposed architecture, scenario figures, and asserted outcomes. Table IV records how those claims are handled here. The purpose is not to discard the idea; it is to keep the research claims proportional to the evidence currently available.

TABLE IV
EVIDENCE AUDIT OF CLAIMS IN THE SUPPLIED DRAFTS

Earlier claim or element	Evidence available in supplied files	Treatment here	Reason
CNN-BiLSTM-attention classifier	Architecture described, but no dataset, weights, training script, splits, or event labels	Retained as candidate	Formula is shown; no accuracy or latency is claimed. Simpler baselines remain necessary.
MAPPO for decentralized control	Proposed in the manuscript; no policy, training environment, reward history, or reproducible results	Optional research candidate	Deterministic constrained optimization and rule-based control form the baseline.
33-kV Ikeja feeder with 42 buses, four ties, eight sectionalizers, and 12 transformers	One draft describes these counts; no circuit file, full branch data, source of the model, or model validation is supplied	Not represented as a validated case	Counts cannot substitute for a complete and independently checked network model.
Fault classification in 11.4 ms and restoration in 1.18 s	Reported in a prior draft; no raw outputs or end-to-end timestamp log supplied	Not claimed as results	The method and the restoration timing cannot be independently reproduced.
SAIDI reduction of 68.4%, SAIFI reduction of 55.1%, and losses falling from 28.4% to 11.2%	Tables in an earlier draft; no source population, loss definition, baseline, code, or matched experiment supplied	Removed from results; listed as targets to test only if explicitly preregistered	These values could otherwise be mistaken for established utility outcomes.
\$480 million CAPEX and 3.8-year payback	Prior draft contains a cumulative investment and savings table but no itemized quantities, exchange-rate date, discounting, asset life, tax, or consistent cash-flow model	Excluded from numerical conclusions	A payback number is not defensible without a cost basis and cash-flow arithmetic.
Failure frequencies and repair-duration ranges for Nigerian urban feeders	Draft table gives ranges without a traceable utility data source	Not reused as Nigerian event rates	Published Lagos outage study is cited separately and limited to its own analyzed records.
Hybrid LoRaWAN/NB-IoT/satellite communication and robotics	Discussed as options, without a site survey, radio-budget calculation, procurement basis, or operational need study	Treated as options for later assessment	Radio choice depends on coverage, latency, energy, maintenance, security, and cost at each site.
Digital twin and federated learning	Reviewed in traceable literature, but not implemented in the supplied project evidence	Included as optional supervisory functions	They are not described as developed software or completed training.

This reconciliation changes the type of contribution. Instead of presenting unsupported output values as if they were experimental results, this manuscript provides a traceable requirements model, equations, controller logic, and a protocol by which the values can later be established. An experiment may show that some draft hypotheses are wrong. That outcome remains useful because the proposed protocol measures feasibility and failure conditions as well as gains.

IX. CYBERSECURITY, SAFETY, AND DEGRADED OPERATION

Distribution automation connects operational technology to measurement, communications, engineering workstations, and in some cases enterprise systems. NIST SP 800-82 Rev. 3 treats OT security in light of physical safety and reliability requirements [16]. The design should apply least-privilege identities, role-based authorization, authenticated command channels, network segmentation, signed or integrity-protected configuration packages, time-synchronized event logs, controlled remote access, and configuration backups. These are proposed controls whose implementation must be checked against the utility's existing architecture and procurement constraints.

The threat model includes unauthorized switching commands, compromised engineering credentials, replayed or delayed messages, spoofed switch status, malicious or accidental configuration changes, corrupted classifier inputs, model-version substitution, and loss of supervisory communications. Controls should include independent confirmation of critical states, command identifiers and sequence numbers, rejection

of expired commands, allow-listed switch operations, two-person or dispatcher authorization for exceptional plans, and audit logs stored outside the single device being controlled. If a cybersecurity alarm affects confidence in command or telemetry integrity, automatic switching should be blocked until an authorized recovery procedure restores trust.

A fault-tolerant design does not mean that the network must continue autonomous operation under every failure. Its safe behavior may be to stop remote switching while local relays continue to protect equipment and staff. The controller must not remotely reclose into a feeder under a lockout, an active work permit, an uncertain earth state, or a physically unconfirmed isolation boundary. These conditions depend on utility procedures and must be implemented as explicit interlocks rather than natural-language warnings. The HSE Code and Distribution Code published by NERC provide the regulatory material against which the utility would review the final implementation [14], [15].

The cybersecurity verification plan includes negative tests: an unauthenticated command must be rejected; an expired command must be rejected; a configuration change outside the approved deployment process must be detectable; a stale switch state must block a dependent command; a lost communication link must not disable local relay operation; and a model service outage must not remove manual control. Test results should include the attempted command, expected behavior, observed behavior, logs, and build version. Compliance cannot be inferred merely from the presence of encryption or a vendor security certificate.

X. ECONOMIC ASSESSMENT WITHOUT UNSUPPORTED COST CLAIMS

The economic model must be built from the proposed pilot's actual bill of materials and utility-approved assumptions. The prior draft's national-scale cost totals and payback figure are not used because the supplied material does not show an auditable component inventory, installation quantity, procurement quotation, operational cost, discount rate, benefit attribution, or consistent yearly cash-flow sequence. The method below states the necessary calculations without filling the missing quantities with guesses.

For a project horizon of H years, initial capital cost C_0 , annual investment C_t , operating cost O_t , measurable benefit B_t , and discount rate r , the net present value is

$$\text{NPV} = -C_0 + \sum_{t=1}^H \frac{B_t - O_t - C_t}{(1+r)^t}. \quad (27)$$

Simple payback is the first year T for which cumulative undiscounted net cash flow is nonnegative:

$$T = \min \left\{ n : -C_0 + \sum_{t=1}^n (B_t - O_t - C_t) \geq 0 \right\}. \quad (28)$$

Discounted payback uses the discounted cash flows in (27). Benefits may include reduced crew dispatch for verified events, lower energy not served, fewer repeated outage visits, and measured technical-loss reduction. Each benefit must be linked to an observable baseline and must avoid double counting. A reduction in customer interruption time is not automatically equal to revenue collected; tariff, billing, collection, and energy-supply conditions influence the financial conversion.

The first financial study should compare three explicitly bounded cases: existing manual processes with the same protection scheme; a limited automation upgrade using existing SCADA and switching equipment; and the proposed supervisory system. Component-level CAPEX should identify sensors, communication, switch controllers, integration, cybersecurity, engineering, commissioning, training, and spare parts. OPEX should include connectivity, licence or maintenance agreements, battery replacement, device calibration, software support, re-training, and periodic security tests. Sensitivity analysis should vary event rates, costs, service availability, communication charges, useful life, discount rate, and achieved benefits. Until these inputs are available, the economic result is a calculation template rather than an estimate of utility payback.

XI. DEPLOYMENT AND RESEARCH STAGES

A staged approach reduces safety and data risk. Stage 1 documents the feeder, protection settings, switching procedures, asset records, existing telemetered points, communications coverage, and event-record quality. Stage 2 builds and independently checks a distribution model and runs baseline power flow and protection studies. Stage 3 implements an offline event replay and advisory tool, which produces recommendations but cannot command switches. Stage 4 tests the controller in a digital environment with injected failures, state corruption, work blocks, and communication degradation. Stage 5 uses hardware-in-the-loop or a switch-controller test bench to verify command




Fig. 3. Proposed evidence-gated progression. A stage advances only after the prior stage's model, security, electrical, and operational acceptance tests pass.

and acknowledgement behavior. Stage 6, if the utility authorizes it, permits a limited field trial under a named engineer, approved test plan, and explicit rollback procedure.

Each stage requires a stop condition. Examples include unverified source model, unresolved switch-position errors, relay-setting conflicts, unsafe action proposals, unacceptable false-positive rate, loss of local protection during communications failure, and inability to reconstruct an event log. The utility may reject a feature or retain manual approval indefinitely. The research outcome does not depend on reaching full autonomy.

The implementation plan also identifies responsibilities. Distribution planning owns the asset and connectivity model; protection engineers own approved settings and coordination studies; SCADA/telecommunications staff own telemetry and command paths; HSE personnel own permit and safety interlocks; operations staff define operating envelopes and control modes; IT/OT security staff review identities, networks, logging, and recovery; and research staff maintain the model, experiments, and evaluation scripts. These responsibilities should be reflected in an approved responsibility matrix rather than assumed from the software architecture.

XII. DISCUSSION

The primary engineering argument is that useful automation is bounded by the quality of state evidence and the utility's authority model. A neural classifier may reduce the time needed to form a fault hypothesis, but it cannot confirm a switch position, determine whether a crew has earthed a conductor, or replace an approved protection setting. Similarly, a reconfiguration algorithm may find a low-loss topology in a simplified model but produce an invalid action when the real feeder has a phase-specific load, a failed switch, a stale measurement, or a work block. The proposed design treats those facts as constraints instead of issues to discuss after optimization.

The EEDC/MainPower training report also changes how the objective should be interpreted. A restoration plan is not complete when a graph says that a node is connected to a source. Safe isolation, load and fuse compatibility, equipment condition, permit status, available material, and the subsequent field repair all affect the operating outcome. For this reason, the

event manager links the electrical sequence to asset and work records, while keeping physical repair confirmation separate from inferred diagnosis. The stores and maintenance records should not be used as proxies for live electrical state; they answer different questions and need their own provenance.

The architecture deliberately postpones the strongest autonomy claims. A rule-based and constrained optimizer can provide a transparent baseline against which neural classification or PPO is tested. If an ML method does not improve end-to-end outcomes under the same constraint checks, it need not be included in the production design. If the data are too sparse for supervised learning, engineers may begin with relay targets, sequence components, waveform statistics, and expert-reviewed event rules. The learning method should be selected after measuring the number, range, quality, and labelling cost of available fault records.

There are also transfer limits. Lagos outage records cannot be treated as a substitute for Enugu or MainPower logs. A standard benchmark demonstrates that a method can be tested on known circuit data; it does not establish that the feeder, customers, asset condition, and operational rules in that benchmark represent a Nigerian utility. Communication coverage, protection settings, device capabilities, and staffing practices vary. An evaluation should therefore report the benchmark, the local calibration data, the conditions that were not modeled, and any change between simulation and field operation.

XIII. LIMITATIONS AND THREATS TO VALIDITY

The present study is a design proposal. The SIWES report supplies practitioner context, but it does not include a representative outage event set, full feeder topology, or utility measurements sufficient to calibrate the equations. The literature synthesis is focused rather than systematic. The proposed control objective uses weights and constraints that require utility approval. The classifier and policy learner are mathematical candidates only; no training or performance result is asserted. The cost model is not populated because site quantities and financial assumptions are absent. These limitations prohibit claims of observed restoration-time or reliability improvement.

Simulation results, if later obtained, will have their own threats to validity. A simplified feeder may omit regulator behavior, phase imbalance, distributed-generation controls, communication delay, and protection interaction. Synthetic fault windows may not reflect real transients or sensor saturation. A test set drawn from the same simulator and parameter ranges as training can overstate generalization. Aggregate means can hide unsafe tail events. A digital twin can also appear accurate if it is repeatedly tuned against the same operating data used to report its performance. Validation should therefore use a held-out feeder or time period, include out-of-range conditions, and preserve a record of all failed cases.

A field trial adds human-factors and operational constraints. Dispatchers may reasonably decline an algorithm's recommendation; that decision is not a software failure. Response time must distinguish the algorithm's calculation time from acknowledgement delays, operator approval, switching-device travel,

and physical repair. Reliability measures need the utility's customer-level event accounting and approved interruption-inclusion rules. The trial should record operator overrides, blocked actions, false alarms, maintenance interventions, communication failures, and any safety-relevant deviation. The evaluation committee must establish stop criteria before commissioning.

XIV. CONCLUSION

This paper proposes a human-on-the-loop cyber-physical architecture for fault-aware restoration in Nigerian distribution feeders. The design starts with requirements grounded in the author's documented EEDC/MainPower industrial training: safe field access, protection coordination, asset and load data, measurement quality, power quality, device condition, material availability, and dependable communications. It combines local protection, edge event analysis, topology and power-flow validation, constrained reconfiguration, operator authority, event logging, and degraded modes. Mathematical formulations define the optimization objective and the conditions under which a switching proposal may be accepted.

The work does not claim a completed simulation or field deployment. Previously reported timing, reliability, loss, and payback values are excluded because the supplied files do not contain the model, data, source code, or cash-flow basis needed to reproduce them. The next research step is to establish a fully documented benchmark experiment, compare rule-based and constrained optimization baselines, inject electrical and cyber failures, and publish the data and scripts needed to reproduce each result. Utility-specific calibration and a limited field pilot should proceed only after engineering review and formal approval.

ACKNOWLEDGMENT

The author acknowledges the engineering staff and supervisors who provided supervised industrial exposure at EEDC/MainPower and the FRCN/Tait Radio Academy programme. The training report is the source of the personal field context; no utility dataset, engineering authorization, funding award, or field trial is claimed in this paper.

APPENDIX A

NOTATION AND MINIMUM INPUT REGISTER

Table V defines the principal variables. A complete model should also store units, source, date, validity period, and owner for each input. Values copied from drawings or test reports should retain the document revision and approval status. When a value is unknown, the register should mark it as unknown and record the conservative action used by the controller.

APPENDIX B

MINIMUM ARTIFACT CHECKLIST FOR REPRODUCTION

Before a result is described as reproduced, the artifact should contain the feeder model, model source, licence, parameter units, all switches and normally-open ties, transformer ratings, source assumptions, per-phase data when applicable, relay

TABLE V
PRINCIPAL SYMBOLS

Symbol	Meaning
$G = (\mathcal{N}, \mathcal{E})$	Feeder graph, nodes and branches
u_{ij}	Branch switch status, 1 closed
z_i	Node energized status
P_{ij}, Q_{ij}	Branch active and reactive flow
v_i, ℓ_{ij}	Squared voltage and squared current
r_{ij}, x_{ij}	Branch resistance and reactance
E_{NS}	Estimated energy not supplied
D_V	Normalized voltage-deviation cost
N_{sw}	Number of switch-state changes
R_{op}	Approved operational risk penalty
U_{unc}	Data/model uncertainty penalty
d_k, a_k	Data validity flag and observation age
c_f, θ_f	Classifier confidence and validation threshold
B	Worksite and maintenance blocks
$\mathcal{A}_{safe}$	Actions admitted by hard safety checks

settings and their approval status, fault-generation procedure, scenario seeds, software versions, scripts, raw outputs, and summary code. Tables should be generated from machine-readable result files, not typed by hand from screenshots. Each result should link to a run identifier. If a figure is illustrative rather than measured, its caption should say so. A reproduction report should list successful cases, rejected actions, timeouts, infeasible topologies, and model failures.

APPENDIX C OPERATOR ACCEPTANCE QUESTIONS

For a candidate restoration plan, the operator interface should make it possible to verify: the trip source and timestamp; whether the protected segment is isolated; whether switch positions are confirmed; which loads will be restored; whether voltage and equipment limits pass; whether relay coordination remains valid; whether any work permit or lockout applies; whether the communications state is current; who authorized the plan; and how the system will respond if a commanded switch does not acknowledge. If one answer is unknown, the interface should identify the unknown explicitly. The software should not imply that an incomplete check has passed.

REFERENCES

- [1] I. Srivastava, S. Bhat, B. V. S. Vardhan, and N. D. Bokde, "Fault detection, isolation and service restoration in modern power distribution systems: A review," *Energies*, vol. 15, no. 19, p. 7264, 2022.
- [2] S. A. Arefifar, M. S. Alam, and A. Hamadi, "A review on self-healing in modern power distribution systems," *Journal of Modern Power Systems and Clean Energy*, vol. 11, no. 6, pp. 1719–1733, 2023.
- [3] F. Shen, Q. Wu, and Y. Xue, "Review of service restoration for distribution networks," *Journal of Modern Power Systems and Clean Energy*, vol. 8, no. 1, pp. 1–14, 2020.
- [4] O. C. Nnaemeka, "A technical report on student industrial work experience scheme (siwes) at enugu electricity distribution company/mainpower and federal radio corporation of nigeria in partnership with tait radio academy," Department of Electronic Engineering, Federal University of Technology, Owerri, Tech. Rep., Nov. 2025, unpublished industrial training report supplied by the author.
- [5] S. Onodjohwo, B. Adebajani, F. Ariba, E. T. Fasina, I. O. Yusuf, and A. Josephine, "Need for network reconfiguration in nigerian distribution systems: A review," in *2024 International Conference on Science, Engineering and Business for Driving Sustainable Development Goals (SEBASDGS)*, 2024.
- [6] S. A. Akanni, A. O. Melodi, and T. O. Ale, "Analysis of the causes of outages in the distribution system of lagos state, nigeria," *FUOYE Journal of Engineering and Technology*, vol. 8, no. 1, pp. 1–8, 2023.
- [7] R. O. Nwajuonye, I. I. Okonkwo, and J. I. Iloh, "Distribution network reliability and efficiency enhancement using distributed generation," *Journal of Engineering Research and Reports*, vol. 21, no. 3, pp. 55–71, 2021.
- [8] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401–1407, 1989.
- [9] IEEE Power and Energy Society, "Ieee guide for electric power distribution reliability indices," IEEE Std 1366-2022, 2022. [Online]. Available: <https://standards.ieee.org/ieee/1366/7243/>
- [10] M. Al-Saadi, M. Al-Greer, and M. Short, "Reinforcement learning-based intelligent control strategies for optimal power management in advanced power distribution systems: A survey," *Energies*, vol. 16, no. 4, p. 1608, 2023.
- [11] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017. [Online]. Available: <https://arxiv.org/abs/1707.06347>
- [12] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agueray Arcas, "Communication-efficient learning of deep networks from decentralized data," *Proceedings of Machine Learning Research*, vol. 54, pp. 1273–1282, 2017. [Online]. Available: <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [13] M. M. H. Sifat, S. M. Choudhury, S. K. Das, M. H. Ahamed, S. M. Mueen, M. M. Hasan, M. F. Ali, Z. Tasneem, M. M. Islam, M. R. Islam, M. F. R. Badal, S. H. Abhi, S. K. Sarker, and P. Das, "Towards electric digital twin grid: Technology and framework review," *Energy and AI*, vol. 11, p. 100213, 2023.
- [14] Nigerian Electricity Regulatory Commission, "The distribution code," NERC code, published July 22, 2018, 2018. [Online]. Available: <https://nerc.gov.ng/resources/the-distribution-code/>
- [15] —, "Health and safety code for the nigerian electricity supply industry," 2026, second edition, March 2026. [Online]. Available: <https://nerc.gov.ng/resource-category/codes/>
- [16] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (ot) security," National Institute of Standards and Technology, Tech. Rep. NIST Special Publication 800-82 Rev. 3, 2023.
- [17] Electric Power Research Institute, "Introduction to opendss," OpenDSS documentation, accessed: October 10, 2026. [Online]. Available: <https://opendss.epri.com/IntroductiontoOpenDSS.html>