



*An Online PDH Course
brought to you by
CEDengineering.com*

AI In Protection, Control, and Cybersecurity for Modern Power Systems

Course No: E01-016

Credit: 1 PDH

Brian Lisiewski, P.E.



Continuing Education and Development, Inc.

P: (877) 322-5800

info@cedengineering.com

Table of Contents

Module 1: Introduction—Safety, Reliability & the Engineer's Role in AI Integration	1
1.1 The Engineer in Responsible Charge.....	1
1.2 Safety and Reliability Culture with AI	1
1.3 Ethics and Policy Alignment	2
Module 2: Protection & Control Fundamentals—Principles and Challenges.....	3
2.1 Modern Grid Stressors—Why Consider AI.....	3
2.2 Trust Boundaries and Safe Operation	4
2.3 Multi-Disciplinary Aspects	4
Module 3: AI-Assisted Protection—Relay Setting Reviews & Disturbance Classification	5
3.1 AI-Driven Relay Setting Analysis	5
3.2 Disturbance and Fault Classification	5
3.3 Benefits and Requirements	6
Module 4: Adaptive Protection & Control—AI's Potential and Limits.....	7
4.1 Use Cases for Adaptive Protection	7
4.2 Boundaries and Safeguards for Adaptivity	7
4.3 Human Oversight of Adaptive Behavior	8
4.4 Example—Microgrid Adaptive Protection	8
Module 5: Cybersecurity for AI-Enabled Protection & Control	9
5.1 Secure Deployment of AI in OT Environments.....	9
5.2 Adversarial ML Risks	9
5.3 Mitigation Strategies.....	10
5.4 Standards and Frameworks	10
Module 6: Model Risk Management, Testing & Reliability Validation.....	11
6.1 Benchmarking AI Performance	11
6.2 Testing with Standard and Worst-Case Scenarios.....	11
6.3 Reliability Indices and Operational Impact	12
6.4 Model Change Management and Continuous Improvement	12
Module 7: Governance, Change Control & Professional Defensibility.....	13
7.1 Policy: The Engineer-in-Loop Mandate	13
7.2 Configuration and Change Management	13

7.3 Dealing with Standards and Regulators	13
7.4 Professionalism and Ethics	14
References	15

Module 1: Introduction—Safety, Reliability & the Engineer's Role in AI Integration

Power system protection and control are safety-critical domains, historically governed by deterministic engineering principles. This module establishes that any introduction of AI into protection/control must augment human decision-making—never replace it.

In these systems, a single misoperation can cause equipment damage, widespread blackouts, or even endanger lives. **AI/ML offers new tools** to analyze complex data and detect subtle patterns, but it cannot compromise the non-negotiable requirements of traditional protection (dependability, security, sensitivity, selectivity) or the responsibility of licensed engineers.

1.1 The Engineer in Responsible Charge

The **Professional Engineer (PE)** is always in charge of protection system performance. **AI does not change the engineer's accountability**: any settings or control actions must still be scrutinized and approved by the PE. For example, if an AI tool suggests altering relay settings, an engineer must verify those new settings with offline studies and ensure they meet all criteria (fault coverage, coordination, margins) **before** they are applied. The NSPE Code of Ethics and NERC standards require that no protective functions be delegated entirely to unverified automated processes without oversight. In short: **AI is like an expert advisor, not an autonomous operator**. The engineer is the one who signs off and ultimately holds responsibility for system outcomes. This ensures that public safety and grid stability are not gambled on an unvalidated black box.

1.2 Safety and Reliability Culture with AI

The goal in adopting AI is to **improve reliability while preserving safety margins**. For instance, AI might detect a fault and isolate it far faster than a conventional relay by analyzing high-speed sensor data, reducing wildfire risk and outage impact. Research prototypes have demonstrated fault isolation dramatically faster in specific wildfire-prevention applications; however, such results apply to those particular systems and conditions and should not be generalized across all AI relay implementations without independent validation.

While such capabilities are promising, they must be applied within a framework where fail-safe defaults still operate. If an AI system fails or behaves unexpectedly, the backup (traditional relay logic or human operator) must catch the scenario.

1.3 Ethics and Policy Alignment

Utilities and regulators maintain that any new automation aligns with established standards. NERC's **Critical Infrastructure Protection (CIP)** guidelines call for maintaining reliability, resilience, and security as the grid digitalizes. Therefore, deploying AI in field devices or control centers requires **thorough documentation** (model and data specifications, test results) and **controlled change management** (just like any protection setting or SCADA scheme change). This allows AI-assisted processes to be auditable and transparent for compliance and engineering review, and ensures that adoption of AI does not outpace the proper development of new standards or training for safe usage.

**AI is an assistant**, not an **authority**

Even as AI speeds up analysis, the PE must **retain final say**. AI can suggest optimized relay settings or identify anomalies, but a licensed engineer must vet these suggestions and ensure they meet all safety and regulatory standards before implementation.

**Faster detection, smarter action**

AI can greatly accelerate fault detection and **isolation**. For instance, DOE & SNL's AI-based relaying can isolate faults **~100× faster** than traditional relays. Faster clearing means improved grid resilience, but only if robustly validated and safely integrated.

Module 2: Protection & Control Fundamentals—Principles and Challenges

Before applying AI, engineers must reaffirm the bedrock of protection and control. This module revisits the fundamentals of power system protection—the objectives and constraints that any advanced technique must uphold.

The complexity of modern systems (renewables, power electronics, demand variability) is straining traditional methods, but any new solution must remain consistent with core protection principles:

- **Dependability:** *Will it trip when it should?* (Never hesitate to trip correctly.)
- **Security:** *Will it restrain when it should?* (Never trip incorrectly.)
- **Sensitivity:** *Will it sense weak or low-magnitude faults?*
- **Selectivity:** *Will it isolate only the faulty section?* (Minimize unnecessary outages.)

These four criteria are the yardstick for any protection scheme. AI cannot compromise them—if anything, AI must be used to **enhance those attributes**. For example, an ML algorithm might boost sensitivity by recognizing small patterns of partial discharges indicative of an impending cable failure, thereby improving dependability. However, in doing so it must maintain security (i.e., avoid false trips due to noise or spurious patterns). As we integrate AI, we will repeatedly refer to these principles to gauge our solution's acceptability.

2.1 Modern Grid Stressors—Why Consider AI

The power grid has become more complex with **inverter-based resources** (solar, wind), **active distribution networks**, and **bidirectional flows**. This creates new challenges for protection:

- **Variability in fault currents:** Inverter-based resources may supply limited or highly controlled fault current, making it hard for conventional relays to detect faults (diminished sensitivity). In some cases they may feed current from multiple directions into a fault, complicating selectivity.
- **Frequent topology and state changes:** Distributed generation, dynamic network reconfigurations, and microgrid islanding mean the system can have multiple operating modes that static settings may not cover optimally.
- **Massive data and events:** Digital substation devices produce large volumes of data (point-on-wave sampling, digital fault recorder files, PMU streams) that are beyond manual analysis scale. AI can help extract actionable patterns from big data.

2.2 Trust Boundaries and Safe Operation

Protective relays are the **last line of defense** for the grid, expected to work flawlessly in microseconds under extreme conditions. This conservative domain has been cautious about trusting new technology (even the shift from electromechanical to numerical relays took decades of proving). Therefore, **any AI integration must fit within existing protection architectures as a complement, not a replacement, and must be fail-safe**. An AI-based advisory system might suggest modifications to settings or highlight miscoordinations, but the actual protective trip decisions still rely on hard-coded logic that meets standards like the IEEE C37 series, unless and until any adaptive logic is thoroughly tested and proven.

2.3 Multi-Disciplinary Aspects

While focusing on electrical fundamentals, we note the interplay with other fields:

- **Mechanical:** Protective controls guard mechanical assets (turbines, transformers) from damage under fault stresses. AI might incorporate mechanical sensor data (vibration, temperature) to anticipate equipment issues, but the protective action still follows electrical fault detection logic.
- **Civil/Structural and Environmental:** Faults can have physical consequences (downed lines causing wildfires). AI applied to protective relaying carries high public safety stakes—e.g., an AI could detect an arcing downed line and rapidly isolate it to prevent igniting a wildfire. Conversely, if misused, AI or communication vulnerabilities could be exploited by attackers to cause outages, a cyber-physical risk addressed in Module 4.

Key takeaway: The bedrock of safe, reliable protection—dependability, security, sensitivity, selectivity—does not change with AI. All stakeholders must insist that any AI-driven approach is assessed against these fundamentals before adoption.

Module 3: AI-Assisted Protection—Relay Setting Reviews & Disturbance Classification

AI algorithms are adept at pattern recognition and big data analysis, making them valuable for supporting protection engineers in complex tasks. This module examines how AI can help with relay setting review and optimization, and fault and disturbance classification.

3.1 AI-Driven Relay Setting Analysis

Setting protective relays remains a meticulous engineering task requiring the balancing of numerous conditions and edge cases. In networks with high renewable penetration or frequent changes, even the best pre-calculated settings may not be ideal for all scenarios. AI can assist by:

- **Automating Coordination Studies:** Expert systems can cross-check and refine settings by comparing them against thousands of contingency scenarios and known coordination criteria. This can help catch potential miscoordination or non-compliance with utility standards.
- **Optimal Setting Suggestions:** Researchers have applied genetic algorithms and fuzzy logic to optimize setting parameters for different system conditions. Any dynamic or "optimal" setting from AI must be validated offline through fault studies before adoption.
- **Online Setting Checking:** Some utilities use AI for real-time monitoring of whether current settings remain secure and dependable under the present system state. If topology or flow changes, an AI engine can simulate fault studies and alert engineers to potential miscoordination.

3.2 Disturbance and Fault Classification

When a protection system operates or an electrical disturbance occurs, post-event analysis is needed to confirm correctness and pinpoint root causes. Machine learning can automate disturbance analysis:

- **Fault Type Identification:** Neural networks and decision trees can identify fault type and location from relay data or waveform snapshots—e.g., classifying whether a fault is phase-to-ground or phase-to-phase based on current transient patterns.
- **Power Quality Event Analysis:** On distribution feeders, AI can classify events such as capacitor switching transients, motor starting, or incipient equipment problems like arcing connections, helping engineers prioritize which events need action.
- **Misoperation Detection:** By comparing actual relay actions to expected responses under similar conditions, ML can highlight anomalies that might indicate a relay misoperation or configuration error.

3.3 Benefits and Requirements

AI-driven analysis can process far more data than humans, potentially increasing coverage and consistency. However, it introduces new governance requirements:

- **Thorough Training and Testing:** Models must be trained on diverse, representative event data and validated to ensure they do not misclassify novel events.
- **False Alarm Control:** It is essential to calibrate thresholds to avoid overwhelming operators with false positives while not missing critical events—a classic precision-versus-recall trade-off.
- **Integration and Transparency:** Advice from AI tools should be presented with explanations (e.g., "possible miscoordination due to CT saturation") so that engineers can trust and verify suggestions.

Key takeaway: AI is a valuable assistant for protection engineers—crunching data and highlighting patterns. It can expedite analysis and propose optimized solutions. However, the human engineer must always validate and contextualize these suggestions before any result is applied without careful review and field testing.

Module 4: Adaptive Protection & Control—AI's Potential and Limits

Adaptive protection refers to adjusting protection settings or logic in response to changing grid conditions. AI can facilitate adaptive protection by rapidly analyzing system state and advising on setting adjustments. However, adaptive protection must be approached cautiously.

4.1 Use Cases for Adaptive Protection

- **DER-Heavy Feeders:** On a feeder with large solar PV, fault current availability and direction vary by time of day. AI can analyze system state in real time (which breakers are closed, generation output, loads) and recommend or switch to appropriate setting groups for different conditions.
- **Wide-Area Adaptive Schemes:** In an integrated grid, line flows and fault levels change with topology. AI can process SCADA topology data and recalculate protection coordination points, ensuring local and remote relay coordination as the system changes.

4.2 Boundaries and Safeguards for Adaptivity

With power system protection, one must be wary of unknown unknowns. Fully autonomous adaptive relaying—completely adjusting settings without human oversight—is generally not acceptable today, except in very limited, pre-tested scenarios:

- **Coordination Complexity:** Interactions among multiple relays could produce unexpected outcomes if changes are not carefully orchestrated. An AI inadvertently reducing a pickup for sensitivity could cause miscoordination with an upstream device, risking widespread outage.
- **Hidden Failure Risk:** Traditional protection has known failure modes. An adaptive scheme that makes an undetected error could introduce a new hidden failure that only reveals itself during a critical fault. Robust failsafe strategies are required.
- **Communication and Latency:** Many adaptive schemes rely on communications. Loss of comms can be as dangerous as a cyber compromise: it might freeze an old setting or mis-synchronize which group is active. A fallback to conservative fixed settings on comms failure is essential.

4.3 Human Oversight of Adaptive Behavior

The industry consensus is that **engineers must supervise adaptive schemes until they achieve the trust equivalent to conventional systems**. In practice:

- If the system is adjusting relay parameters in real time (e.g., switching a pre-calculated setting group on islanding), the logic for that switching is usually hardwired and triggered by an authoritative signal—not an opaque ML decision. Each such scheme is carefully studied, validated with staged tests, and the risk of malfunction is weighed before enabling.
- More advanced AI-driven adaptation is currently confined to research and pilot trials—and if used, typically in advisory mode. A protection engineer reviews suggestions, tests them offline, and explicitly approves and loads them during a scheduled maintenance window.

4.4 Example—Microgrid Adaptive Protection

Consider a microgrid that can grid-connect or island. A multi-agent system with AI logic might coordinate overcurrent relays adjusting their settings in real time depending on the mode and PV output. While demonstrated in labs to improve reliability, this raises a critical question: *what if an attacker hacks the communication among these intelligent agents?* False signals could cause a failure to trip for a real fault or a false trip for a non-fault—turning adaptivity into a vulnerability. Any foray into highly adaptive protection must go hand-in-hand with robust cybersecurity (Module 5).

Key takeaway: Adaptive protection, enabled by AI, is powerful but double-edged. It must be restrictively bounded by what is proven safe, with careful design to maintain traditional reliability principles. Most adaptive schemes today are limited to switching between pre-validated setting groups or providing recommendations that an engineer verifies before adoption.

Module 5: Cybersecurity for AI-Enabled Protection & Control

Embedding AI in power system operations introduces new cybersecurity considerations. Not only do we need to protect the power system with AI; we must also protect the AI itself and the underlying OT systems from cyber threats.

5.1 Secure Deployment of AI in OT Environments

Power system protection and control devices—such as microprocessor relays, RTUs, and PLCs—typically operate in constrained and highly secure environments. Introducing AI often means adding new software or connections which can **expand the attack surface** if not managed properly. Best practices include:

- **Network Segmentation and Hardening:** AI processes should abide by the same segmentation rules as other critical OT components (e.g., placed in secure zones with minimal external connectivity, or isolated by data diode where only one-way data flow is required). Connections to cloud resources or vendor remote access must be carefully examined and likely disallowed or confined to test environments, given CIP-005's strict limitations on remote access for critical systems.
- **Access Control and Authentication:** Only authorized personnel should be able to modify AI model parameters or feed it new training data. Use multi-factor authentication and strict role-based access control for any interfaces, consistent with CIP-007 guidance.
- **System and Data Integrity:** Ensure the pipeline feeding sensor data to AI models is protected by encryption and authentication per IEC 62351 for secure SCADA communications. If an attacker alters data inputs, they could fool the model—a risk known as false data injection (FDI). Use secure protocols (e.g., TLS for communications, code signing for model files) and monitor for anomalies that might indicate tampering.

5.2 Adversarial ML Risks

Beyond standard OT security, **adversarial machine learning** is an emerging concern. These are tactics by which an attacker exploits the ML model itself:

- **Feeding Malicious Inputs:** Attackers could craft inputs that cause an AI model to misclassify events or make bad recommendations—e.g., injecting noise that confuses an anomaly detector, suppressing an alarm, or causing a false trip. Research has demonstrated the potential for such data poisoning and spoofing in control systems.
- **Model Tampering:** If an attacker gains access to the platform running an AI model, they might alter the model to degrade protection—for example, maliciously retraining it to under-report everything as normal (compromising dependability) or to generate constant false alarms (overwhelming operators).

5.3 Mitigation Strategies

Incorporate adversarial resilience into the design:

- Use robust ML techniques less sensitive to out-of-distribution inputs, or with built-in checks. For example, use ensemble models requiring consensus for high-confidence protective actions, or restrict the domain of AI suggestions to within known safe bounds.
- Never allow an ML model to directly execute control actions without a sanity check by conventional logic or a human. If an AI suggests tripping a line, ensure a conventional protection element or operator confirmation concurs before issuing the trip.
- **Continuous Monitoring:** Treat the AI system itself as a critical asset to monitor. Use anomaly detection on the model's behavior—e.g., if it suddenly starts producing an unusually high volume of high-criticality alarms, that could indicate either a fault condition or an attack.

5.4 Standards and Frameworks

Regulators and industry groups are beginning to address these new issues:

- **NERC CIP:** Standards CIP-005 (Electronic Security Perimeters), CIP-007 (Systems Security Management), and CIP-010 (Configuration Change Management) require baseline configurations and change management for critical systems. An AI model integrated into protective relaying would likely be considered part of the BES Cyber System and thus subject to CIP change controls: no unauthorized modifications, rigorous testing of changes, and documented back-out plans.
- **Security by Design:** Following NIST SP 800-82 (Guide to ICS Security) and IEC 62443 (Security for Industrial Automation and Control Systems) ensures that adding AI does not inadvertently weaken security. Any new data link for gathering PMU data should be protected as if it were a control signal, with encryption and intrusion monitoring.

Key takeaway: The integration of AI into protection and control requires a dual focus: improving the physical reliability of the grid and not introducing cyber vulnerabilities. Public safety and system integrity extend beyond physics to cybersecurity—a compromised AI or data pipeline can be as damaging as any physical fault.

Module 6: Model Risk Management, Testing & Reliability Validation

The guiding principle when applying AI in protection is to test rigorously and verify all outputs before trusting any recommendation in the field. — or simply drop the quotation marks: Trust but verify is the guiding principle when applying AI in protection.

6.1 Benchmarking AI Performance

To gauge an AI model's readiness, it must be tested on known scenarios and measured with quantitative metrics:

- **Accuracy/Precision/Recall:** For classification tasks, measure how well the model identifies true events versus false ones. High precision ensures few false alarms (security); high recall ensures catching most actual issues (dependability). For instance, published research reports ML fault classifiers achieving up to 98% accuracy under controlled laboratory conditions; however, field deployment performance can differ significantly due to sensor noise, variation, and novel fault types not present in training data.
- **Mean Error and Bias:** For predictive tasks, track errors (e.g., mean absolute error or bias). If an outage probability model systematically underestimates risk under specific conditions, adjust outputs conservatively or retrain with more representative data.
- **Latency and Throughput:** If the AI runs in real time, ensure it meets required speed. Protective actions often must occur within milliseconds. Confirm the model and its data pipeline operate within time constraints—verify that the full end-to-end system (data acquisition, inference, and action) fits the clearing time requirements.

6.2 Testing with Standard and Worst-Case Scenarios

Adopting AI in protection should involve rigorous offline testing akin to type tests for protective relays:

- **Playback of Fault Records:** Use historical fault and disturbance records (DFR event files, COMTRADE records). Does the AI correctly identify and respond as expected? Feed actual misoperation cases (e.g., from NERC system protection reports) to verify the AI would have caught them, and feed borderline cases to confirm it would not have mis-operated.
- **Standards-Based Tests:** Use guidelines from IEEE and NERC reliability standards. NERC PRC-002 and PRC-004 deal with disturbance monitoring and relay misoperations—an AI's performance can be evaluated on whether it would help meet these requirements.
- **Edge Conditions:** Stress-test the model with extreme scenarios (N-2 contingencies, very rare fault types, unusual system configurations) to verify it does not output something dangerous or uncertain. Incorporate failsafe logic for out-of-scope conditions.

6.3 Reliability Indices and Operational Impact

Success of AI enhancements can be partially measured by improvements in reliability outcomes:

- Monitor relay misoperation rate—how often a relay trips when it should not or fails when it should. If AI analysis leads to corrections that reduce misoperations, that is a tangible benefit.
- Track changes in SAIDI/SAIFI due to predictive and faster-acting measures. A utility pilot using predictive failure analytics might report reduced feeder outages year-over-year and faster restoration contributing to lower SAIDI.
- Ensure that reliability indices are not improved at the cost of unreported near-misses or hidden reliability debt. Decisions must have thorough audit trails.

6.4 Model Change Management and Continuous Improvement

AI models should be subject to a structured change control process:

- If the model is updated (retrained with new data or parameters), treat it like a new protective relaying setting change: require testing, documentation, and a peer review by another engineer before going live. This aligns with ISO 9001's emphasis on controlled processes.
- Maintain a model logbook—record training data versions, model versions, reasons for changes, validation results for each version, and who approved it. If a question arises ("Why did the model miss this event?"), you can trace exactly which version was in use and what its known limitations were.
- Continuously gather feedback: every time a real event occurs, compare model output to reality. Any false negative (missed detection) or false positive (unnecessary alert) is an opportunity to refine the model.

Key takeaway: Ensuring AI models in protection and control do what they are supposed to—and nothing more—requires a robust regime of testing, metrics, and process control. AI models demand the same lifecycle management as protective relays: type testing, acceptance testing, maintenance testing, and periodic setting reviews.

Module 7: Governance, Change Control & Professional Defensibility

In the final module, we synthesize lessons on implementing AI in grid protection and control within a governance framework that maintains accountability and public trust.

7.1 Policy: The Engineer-in-Loop Mandate

Organizations adopting AI in operations often codify that **the technology aids, but does not override, professional judgment**. A utility might require that any AI-driven recommendation (a proposed relay setting change or an outage prediction prompting switching) is reviewed and approved by a protection engineer or system operator. Over time, if trust in a specific narrow AI function grows (through a track record of flawless performance under supervision), some decisions might be semi-automated but still with direct human oversight or an immediate human abort capability in emergencies.

7.2 Configuration and Change Management

- **Documentation:** Document everything about the AI's integration—design intent, model architecture summary, data sources, normal operating range, known failure modes. This document should be maintained and accessible like a relay settings database or a protection philosophy document.
- **Versioning and Approvals:** Use formal change management (a "Model Change Request" process) for any modifications. A model or its threshold should have an owner (likely the protection engineering department), and any change requires review, testing evidence, and sign-off—just as one would not change a transformer differential relay's sensitivity without engineering approval.
- **Audit Trails:** As part of CIP-010 compliance and general quality assurance, maintain logs of when the model was consulted and what it output. If an AI suggests reconfiguration, log that suggestion, the data input, and the operator's action (taken or not taken). This is analogous to sequence-of-events records but for the AI's role in operations.

7.3 Dealing with Standards and Regulators

If a protection misoperation or major event occurs involving an AI component, investigators will ask: did the utility exercise prudent practice in using this technology? A utility should be able to show:

- The AI system was extensively tested prior to use, with test reports available.
- A PE oversaw the integration, with documented approvals.
- The AI's actions or recommendations were consistent with, or did not conflict with, existing standards.
- There was a backup plan: if an AI failure is suspected, the system can automatically revert to conventional protection only.
- The utility followed industry guidance or best practices and can reference alignment with EPRI guidelines, IEEE working group outputs, or similar.

7.4 Professionalism and Ethics

The NSPE Code obliges engineers to approve only designs safe for public welfare. By extension, if AI is used, the engineer's duty is to **thoroughly understand its behavior, limitations, and potential failure modes** before trusting it. One should never deploy an opaque tool in a safety-critical role without being able to explain and justify its outputs. Strategies like simplified surrogate models or constraints that align with known principles can make a model explainable to the engineers using it—fostering trust and ensuring that engineers can defend decisions influenced by the model.

Example: "Our ML-based system predicted an 80% likelihood of a feeder fault under the forecasted wind conditions, so we took the prudent action of pre-emptively sectionalizing part of the feeder. This decision was supported by historical analysis showing similar conditions led to faults, and it aligns with our operating procedures for storm mitigation."

Conclusion: Integrating AI into protection, control, and cybersecurity is a journey, not a destination. It requires blending cutting-edge technology with time-tested engineering discipline. By maintaining rigorous oversight, thorough validation, secure and controlled implementations, and by clearly documenting and communicating all actions, PEs can harness AI's benefits—improved speed, insight, and efficiency—without compromising the fundamental mission of keeping the power system safe, reliable, and secure.

References

1. U.S. Department of Energy, Office of Cybersecurity, Energy Security, and Emergency Response (CESER), "Fortifying America's Electric Grid Against Wildfires with AI," energy.gov, 2022. [Online] Available: <https://www.energy.gov/ceser/articles/fortifying-americas-electric-grid-against-wildfires-ai>
2. U.S. Department of Energy, Office of Scientific and Technical Information, "AI/ML for Power System Protection and Control," OSTI Report 1863312, 2022. [Online] Available: <https://www.osti.gov/servlets/purl/1863312>
3. R. Dubey, A. Popov, S. Kazemi, et al., "AI and Machine Learning in Power System Protection," in *Advances in Power Systems and Energy Management*, Springer, 2022. [Online] Available: https://link.springer.com/content/pdf/10.1007/978-981-19-1528-4_74.pdf
4. H. Afrasiabi, M. Mohammadi, M. Rastegar, et al., "Deep Learning-Based Fault Detection and Classification in Power Systems," *Scientific Reports, Nature*, vol. 15, 2025. [Online] Available: <https://www.nature.com/articles/s41598-025-19634-y>
5. North American Electric Reliability Corporation (NERC), CIP-005-7: Electronic Security Perimeters. [Online] Available: <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>
6. North American Electric Reliability Corporation (NERC), CIP-007-6: Systems Security Management. [Online] Available: <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>
7. North American Electric Reliability Corporation (NERC), CIP-010-4: Configuration Change Management and Vulnerability Assessments. [Online] Available: <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>
8. National Institute of Standards and Technology (NIST), SP 800-82 Rev. 3: Guide to Operational Technology (OT) Security, 2023.
9. International Electrotechnical Commission (IEC), IEC 62443: Security for Industrial Automation and Control Systems.
10. International Electrotechnical Commission (IEC), IEC 62351: Power Systems Management and Associated Information Exchange — Data and Communications Security.
11. IEEE, IEEE C37 Series: IEEE Standards for Protective Relaying. [Online] Available: <https://standards.ieee.org>
12. National Society of Professional Engineers (NSPE), NSPE Code of Ethics for Engineers. [Online] Available: <https://www.nspe.org/resources/ethics/code-ethics>
13. International Organization for Standardization, ISO 9001:2015 — Quality Management Systems — Requirements, 2015.