



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

AI For Grid Operations: Event Detection, Outage Prediction & Power Quality Analytics

Course No: E01-015

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 — AI as Advisory for Safe, Reliable Grid Operations	1
1.1 AI's Promise and Responsibility	1
1.2 Trust and Transparency	1
1.3 Multi-Disciplinary Relevance.....	2
Module 2: Operational Data Streams — SCADA, AMI, PMU and Streaming Analytics..	3
2.1 Key Features for AI Models	3
2.2 Real-Time Streaming versus Batch Analysis.....	4
2.3 Practical Example	4
Module 3: Event and Anomaly Detection — Identifying Problems Early	5
3.1 The Need for Early Detection	5
3.2 Anomaly Detection Approaches	5
3.3 False Alarms versus Missed Detections	5
3.4 Integration into Operations	6
3.5 Example	6
3.6 Partial FMEA: AI-Based Event Detection in Operations	7
Module 4: Outage Prediction and Incipient Failure Detection	8
4.1 Outage Prediction using AI.....	8
4.2 Operationalizing Outage Predictions	9
4.3 Incipient Failure and Degradation Detection.....	9
4.4 Reliability Impact	10
4.5 Example — Transformer Failure Averted.....	10
Module 5: Power Quality Analytics — Classification and Root-Cause Acceleration	11
5.1 Types of PQ Disturbances.....	11
5.2 Root Cause Analysis and Action.....	11
5.3 PQ Monitoring Systems.....	12
5.4 Verification and Regulatory Reporting.....	12
5.5 Example	12
Module 6: Validation Metrics and Reliability Indices — Measuring Success	13
6.1 Model Performance Metrics	13
6.2 Reliability Indices as Outcome Metrics.....	14

6.3 Dealing with Uncertainty	14
6.4 Communicating Uncertainty and Results	15
Module 7: Governance, Model Management and Professional Defensibility	16
7.1 Model Governance and Configuration Control	16
7.2 FMEA Revisited — Model Use Risks	16
7.3 Ensuring Regulatory and Ethical Compliance	17
7.4 Communication and Culture	17
7.5 Professional Defensibility	17
References	18

Module 1: Introduction — AI as Advisory for Safe, Reliable Grid Operations

Modern grid operations generate **massive data streams** (from SCADA sensors, smart meters, PMUs, etc.) that present an opportunity for AI to transform how we manage the grid. **AI/ML algorithms can analyze these data streams in real-time or near-real-time to detect events and predict issues** before they escalate.¹ However, **AI must be used as a support tool** — *not* an auto-pilot replacing human expertise. This module establishes that **Professional Engineers remain accountable** for decisions influenced by AI, and must ensure that any automated insights are **thoroughly validated** and implemented **with caution**.

1.1 AI's Promise and Responsibility

AI can **enhance grid reliability and safety** by anticipating problems rather than just reacting to them. For example, an ML model might predict an elevated outage risk on a feeder 48 hours ahead due to weather and asset condition data,² enabling crews to act proactively. Such capability translates into improved reliability metrics (e.g., one reported benchmark scenario suggests AI-managed grids can reduce outage duration — SAIDI — from approximately 142 minutes/year to approximately 38 minutes/year³; real-world results vary widely depending on grid topology, climate, and implementation maturity). However, **the engineer's responsibility** is to verify these predictions, shape them into safe operational actions, and ensure **no compromise to safety or compliance**. A classic example: if an algorithm recommends reconfiguring feeders to offload a transformer under stress, the PE must confirm the new configuration still satisfies voltage limits, protection coordination, and safety clearances before acting. The **NSPE Code of Ethics** reminds us that an engineer's **paramount duty is to hold public safety above all** — thus any technology (no matter how advanced) must be subordinate to that duty. It is up to the engineer to **review AI outputs and intervene** if something does not pass a "reasonableness" test.

1.2 Trust and Transparency

To use AI effectively in operations, **transparency and traceability** are key. Operators and engineers must understand *why* an AI system is flagging an anomaly or predicting an outage — or at least have confidence in the evidence supporting it. If a model can highlight the salient features (like a spiking transformer temperature or a pattern of small voltage dips) that led to an alarm, it is easier for engineers to trust and act on it. This calls for **explainable AI** and **audit trails**: for each critical decision, there should be logs or details (data snapshots, model confidence scores, etc.) so that a human can later review what the model saw.⁴ Moreover, **ISO 9001 quality management** principles encourage thoroughly documented processes — applying that to AI means documenting model versions, training data, validation results, and decision logs, which becomes part of the operational record and can be audited similarly to how protective relay settings or maintenance logs are.

1.3 Multi-Disciplinary Relevance

The topic spans electrical engineering (grid operations), mechanical (e.g., incipient failures in transformers or rotating machines), civil (weather impacts on overhead lines, structural loading by equipment failures), and environmental (power quality affecting sensitive equipment). Engineers in each area will find AI tools relevant:

- Electrical engineers may use ML to predict substation faults or classify power quality events.
- Mechanical engineers might interpret transformer sensor anomalies as mechanical or thermal stresses that a model flags for early maintenance.
- Civil engineers may input weather patterns (like icing or wind loading on lines) to predict structural failures.
- Environmental engineers might see PQ analytics as ensuring compliance with power quality standards that protect sensitive community equipment.

Insight: The goal is to **transition grid operations from reactive to predictive**, but always with an engineer guiding the transition.⁵ AI can monitor far more data in parallel and faster than humans can, recognizing subtle patterns or early warnings — which, when verified by engineers, leads to timely interventions and improved reliability **without sacrificing safety**. The rest of this course will delve into the specifics of these AI applications and outline best practices to implement them responsibly.

👉 AI complements, engineers decide

AI/ML models can spot grid issues early, but the PE must verify and approve any responses. This collaborative approach boosts reliability while upholding safety and compliance.

⚡ Proactive reliability gains

Data-driven outage predictions let utilities **manage outages proactively**, improving uptime and customer satisfaction. Integrated data hubs (weather, SCADA, GIS, etc.) help foresee and avoid events, trimming emergency costs.

Module 2: Operational Data Streams — SCADA, AMI, PMU and Streaming Analytics

A fundamental step in applying AI to grid operations is understanding and leveraging the wealth of operational data sources available. This module reviews key data streams:

- **SCADA (Supervisory Control and Data Acquisition):** Traditional SCADA systems provide periodic snapshots (every few seconds) of grid parameters like feeder currents, voltages, breaker status, etc. It is a backbone for control rooms, used historically for alarm thresholds (limit violations). For AI, SCADA offers rich time-series signals at substation-level and feeder-level resolution.
- **AMI (Advanced Metering Infrastructure):** Smart meters deliver fine-grained usage and voltage data at customer endpoints (commonly every 15 minutes to hourly). While not real-time like SCADA, **AMI data at scale** can reveal consumption patterns, localized voltage issues, and provide ground-truth for outage detection (e.g., multiple meters dropping out indicates an outage location).⁶
- **PMU (Phasor Measurement Unit) or Synchrophasors:** These high-speed sensors (often at transmission substations, increasingly in distribution) measure phase angles and magnitudes at reporting rates typically ranging from 10 to 120 frames per second (per IEEE Std C37.118), with higher rates used for capturing fast transients. They can capture fast transients and oscillations — valuable for **power quality and stability** analysis. PMU data allows detection of dynamic events (like a sudden oscillatory transient) that slower data might miss.

Other sources include **fault recorders**, **weather sensors**, and **asset condition monitors** (like transformer gas or thermal sensors). **AI streaming analytics** platforms ingest these heterogeneous streams and apply algorithms to find patterns and anomalies in real time. The challenge: these data streams are voluminous and often siloed. A robust AI implementation requires a **data integration layer** (sometimes called a "reliability data hub"⁷) to merge SCADA, AMI, PMU, weather, GIS, etc. into a unified view.

2.1 Key Features for AI Models

- **Temporal context:** Because grid data are time-series, features often include lagged values or moving statistics. E.g., a feeder's one-hour rolling average voltage plus the current value can help detect a slow voltage drift anomaly.
- **Spatial context:** Combining data across neighboring sensors can improve detection. For instance, if adjacent feeders both see a slight voltage dip, it might point to an upstream issue (substation transformer problem) which an algorithm can infer by correlating those SCADA points.
- **Weather and exogenous features:** Integration of weather data (wind, temperature, lightning strike data) is crucial for outage prediction. Approximately 70% of outage events on distribution grids are weather-related, so including, say, forecasted wind gusts as features in an outage risk model significantly improves predictive power.⁸

- **Asset data:** Maintenance records or known health indices of equipment (e.g., a transformer's health score, years in service) are important features for incipient failure models.

2.2 Real-Time Streaming versus Batch Analysis

Some AI applications run in real-time stream processing (for immediate alerting), others run in batch (e.g., overnight outage risk for next day). For example, an anomaly detection algorithm might continuously monitor SCADA and PMU feeds, issuing an alert within milliseconds when an unusual pattern is recognized. In contrast, an outage prediction model might run once daily using the latest weather forecast, updated soil moisture levels (affecting pole strength, etc.), and current grid state, producing a risk ranking of circuits. **Engineers must ensure the integration is fail-safe** — if the AI system goes offline or data is delayed, there should be fallback to normal operations (for example, defaulting to normal SCADA threshold alarms if the anomaly detector is unavailable).

2.3 Practical Example

A utility combines feeder SCADA with AMI "last gasp" messages (meters sending a signal when they lose power) to identify outages faster. An ML model processes these data plus momentary voltage dips to pinpoint fault locations more accurately. Additionally, **high-speed PMU data** feeding into an anomaly detector might catch a subtle oscillation that indicates a failing capacitor bank — something a slow SCADA poll might not notice beyond minor voltage fluctuations. By merging these data sources, the utility can see both **broad patterns** and **localized anomalies** that otherwise would remain hidden, allowing faster and more targeted responses.⁹

In summary, having robust data pipelines from SCADA, AMI, PMUs, etc., and proper feature extraction, is the foundation of any successful AI for grid operations. The next modules will discuss how these data are used in AI models to detect events, predict outages, and analyze PQ — all in service of a safer, more reliable grid.

Module 3: Event and Anomaly Detection — Identifying Problems Early

3.1 The Need for Early Detection

Distribution operators have always relied on alarms (voltage out of bounds, frequency deviations, etc.) to catch problems. However, these threshold-based alarms can be too late or too simplistic, and they may miss subtle precursors to failure. AI-enhanced **anomaly detection** can sift through streams of SCADA/PMU data to spot unusual patterns indicating early stages of trouble — e.g., a transformer beginning to fail might exhibit small but progressive increases in operating temperature or odd switching transients that are below traditional alarm thresholds. Catching these anomalies early allows maintenance to be scheduled or loads to be shifted **before a major outage occurs**, thus directly improving reliability and safety.^{10 11}

3.2 Anomaly Detection Approaches

Common AI techniques for event detection include:

- **Unsupervised learning** (no need for labeled "failure" examples) to detect outliers from normal behavior. For instance, an **autoencoder** can be trained on normal SCADA/PMU data; if new data reconstruct poorly (high error), it signals an anomaly.^{12 13} This can catch novel events as well, but may need tuning to minimize false alarms.
- **Supervised classification** of known event types. If a dataset of labeled events exists (e.g., known flashover events, oscillation events), a classifier (like a neural network or random forest) can be trained to recognize these patterns. For known event types (capacitor switching transients, voltage sags due to faults, etc.), supervised models can achieve **high detection accuracy** — research shows deep learning classifiers can reach approximately 95–99% accuracy on multi-class power quality event datasets in controlled benchmark conditions; field accuracy varies depending on noise levels, sensor calibration, and whether disturbance types encountered were represented in training data.¹⁴ However, supervised classifiers only detect what they were trained on.
- **Hybrid methods**: using physics rules plus ML. For example, a rule might flag if voltage falls more than 10% anywhere, then an ML model processes the waveform shape to decide if it is a temporary motor start versus a fault developing, refining the alarm significance.

3.3 False Alarms versus Missed Detections

A critical part of deploying anomaly detection is managing **false positives** (false alarms) versus **false negatives** (missed events). Too many false alarms can desensitize operators (the "cry wolf" effect), while missing real events undermines the whole purpose. **Risk-based tuning** is needed: for life/safety-critical anomalies, you might accept more false alarms to ensure nothing is missed; for minor power quality blips, you might prefer fewer false alarms so staff are not overwhelmed. Many systems incorporate **confidence scores** or multiple thresholds (e.g., a "warning" level versus a "critical" alarm) to calibrate responses. **Engineers must validate these settings** through pilot testing and gradually adjust to strike the right balance.

For instance, one could run the anomaly detector in shadow mode (not triggering real actions) for several months to measure how often it fires and correlating with real events, then tweak sensitivity before fully integrating it into operations.

3.4 Integration into Operations

Detected anomalies must be tied into the control room workflow or maintenance processes. This might mean the AI system opens a case/ticket automatically in an incident management system with relevant data attached. For example, if a pattern of brief voltage sags on multiple feeders is detected, a case could be created suggesting "**Suspected substation issue — investigate transformer tap changer malfunction**" with supporting data. **Audit trails** should record what data triggered the alert and what the AI rationale was (e.g., "voltage unbalance anomaly started at 14:23 on Feeder A" etc.), so an engineer can later review and either confirm it was a valid detection or label it as a false alarm for model retraining.

3.5 Example

A mid-size utility deployed a deep learning anomaly detector on PMU data at the substation level. In one case, the system identified a subtle harmonic distortion pattern that deviated from normal by only approximately 2% (well below the standard alarm threshold) and persisted for several minutes, prompting an alert. Engineers investigated and found an **incipient cable insulation fault** — partial discharges were injecting harmonics before any protective relay operated. They scheduled a controlled repair, avoiding a potential unplanned outage and possible safety hazard had the cable failed catastrophically. This anecdote illustrates how AI can serve as an **early warning system**, but also how essential it is that a human verify the situation before, say, dispatching a crew or flipping switches.

3.6 Partial FMEA: AI-Based Event Detection in Operations

The table below lists representative failure modes, their effects, and mitigating controls for AI-based event detection. Each utility should tailor a similar risk analysis to their system, aligning with **ISO 9001's risk management** emphasis, ensuring robust processes around AI integration.

Failure Mode	Effect	Severity	Likelihood	Risk	Mitigation/ Control
False alarm flood	Operators desensitized; real events missed	High	Low	Medium	Multi-level alert thresholds (warning/critical); shadow-mode tuning before go-live
Missed critical event	Undetected incipient failure leads to unplanned outage or safety hazard	High	Low	High	Tune recall for safety-critical assets; maintain traditional SCADA alarms as backstop
Model offline / data feed failure	No AI-based alerting; reliance on threshold alarms only	Medium	Medium	Medium	Fail-safe fallback to standard SCADA alarms; uptime monitoring and auto-recovery
Stale training data (model drift)	Model accuracy degrades as grid conditions change (new DER, load patterns)	Medium	Medium	Medium	Scheduled periodic retraining and validation; track performance KPIs over time
Misclassified PQ event	Wrong root-cause direction; wasted crew dispatch or missed mitigation	Medium	Medium	Medium	PE review of high-severity classifications; periodic ground-truth sampling
Unauthorized model change	Untested model deployed in production; unpredictable outputs	High	Low	Medium	Change-management process; version control; management sign-off required for production updates

Module 4: Outage Prediction and Incipient Failure Detection

While anomaly detection catches events in progress or about to occur, outage prediction looks ahead to foresee events hours or days in advance. The ability to anticipate outages or equipment failures allows utilities to convert unplanned outages into planned maintenance or targeted reinforcement — drastically improving reliability.

4.1 Outage Prediction using AI

Outage prediction merges **historical outage data**, **real-time conditions**, and **external factors** (like weather forecasts) to estimate the likelihood and potential location of future outages. Traditional methods included statistical models and heuristic scores (e.g., risk indices based on equipment age and load). AI approaches can handle far more variables and complex interactions:

- **Weather-based Outage Prediction:** ML models (e.g., random forests, gradient boosting) can learn from past storm data which weather parameters (wind speed, precipitation, lightning strikes) are most predictive of outages at specific locations.¹⁵ They can output a risk score per feeder or region for the next 24–72 hours, which operators use to pre-stage crews or reconfigure the network preemptively. One real-world example illustrates the scale of impact: Eversource Energy's AI-based outage model, integrating weather, vegetation, SCADA, GIS, and more, reportedly helped avoid approximately 40,000 customer outages in two months by enabling proactive actions.¹⁶
- **Asset Failure and Incipient Faults:** By monitoring asset condition data (transformer oil gas levels, breaker operation counts, partial discharge sensors) and patterns gleaned from SCADA/PMU, AI can flag equipment likely to fail in the coming days. For instance, if a substation transformer shows a combination of rising temperature, hydrogen gas parts-per-million trends, and recent voltage regulation anomalies, an AI model might predict it has a high probability of failure within 72 hours. Some systems have been able to predict specific equipment failures with approximately 60–90% accuracy for well-monitored assets.¹⁷ With, say, 48 hours notice, a utility can deliberately offload that transformer and schedule an emergency replacement before a failure knocks out customers unexpectedly.

4.2 Operationalizing Outage Predictions

Making predictions actionable involves:


- **Prioritization:** Not all risk predictions warrant immediate action. Utilities set thresholds — e.g., "if outage probability exceeds a set level or risk score is beyond a defined value, then create a preventive work order." Lower risks might just be monitored until the next model update.
- **Integration with Outage Management Systems (OMS):** Predicted outage-prone areas might be flagged in the OMS or Advanced Distribution Management System (ADMS), which can then simulate network reconfigurations or prepare switching sequences to isolate minimal customers if an outage happens.
- **Crew and Inventory Staging:** A very practical use — if a model identifies a particular feeder or region as high risk due to an impending severe storm (like heavy wet snow and high winds), the utility can **pre-position repair crews and necessary equipment** near that area. This shortens response times dramatically, as crews are not waiting for calls but are already in the vicinity of likely damage areas. That translates to lower SAIDI/SAIFI since some outages might be prevented or restored faster.¹⁸

4.3 Incipient Failure and Degradation Detection

Not all failures are weather-induced or sudden. Many are preceded by **tell-tale signs**: e.g., a transformer might show a pattern of small gassing events or incremental impedance changes, a cable might show partial discharge noise on a sensor, a breaker might operate slower in its timing tests. AI algorithms can learn what combinations of these subtle signs reliably precede failures. One approach is **time-series classification** — using historical data from before known failures to train models to recognize those precursors and differentiate them from normal drift. The challenge is typically a lack of numerous failure examples (failures are thankfully relatively rare), but techniques like **transfer learning** or **expert-labeled synthetic data** can help models generalize from limited examples. For critical assets, the threshold for raising an incipient failure alarm can be quite low (preferring to check a false positive than miss a real issue). When an ML model flags such a condition, an engineer can decide to **take immediate actions** like moving load off an at-risk transformer or inspecting it with infrared cameras — akin to preventive medicine for the grid.

4.4 Reliability Impact

All these predictive measures tie back to improving reliability indices. **SAIDI** (duration of outages) drops when we prevent or shorten outages via early warning. One benchmark reported a drop in SAIDI from 142 to 38 minutes/year by adopting predictive measures¹⁹ — utilities should establish their own baselines and track improvements incrementally. **SAIFI** (frequency of outages) also improves since some outages are altogether averted. Over time, these improvements not only benefit customers (fewer and shorter outages, better power quality), but also help utilities meet regulatory reliability targets and avoid penalties. As an engineer overseeing AI-based reliability improvements, one might present these gains in terms of **reduced SAIDI/SAIFI** to regulators or management, with the backing of data (e.g., number of outage incidents prevented, improvement in response times, etc.).



4.5 Example — Transformer Failure Averted

A large industrial park's substation transformer had been in service 25 years. An ML model analyzed its SCADA load data, oil temperature trend, ambient temperature, and previous tap changer operations. It detected an **anomalous heating trend** (temperatures creeping above normal on moderate days) with intermittent micro interruptions (millisecond-range events seen in PMU data), and assigned a high probability of an incipient winding insulation failure within weeks. The utility prioritized a **planned replacement** of that transformer the next weekend. Post-replacement inspection confirmed signs of impending failure (charring, acidic oil). **This prevented a potentially catastrophic mid-summer outage** that could have affected critical industrial processes. The example illustrates how AI predictions, **when acted upon by engineers**, can significantly boost reliability and safety — but it also shows the process: the AI flagged something unusual, engineers validated (cross-checking with field measurements and asset history), then proactively intervened.²⁰

Module 5: Power Quality Analytics — Classification and Root-Cause Acceleration

Power quality (PQ) disturbances — voltage sags, swells, harmonics, flicker, transients — are daily occurrences on power systems. While they may not always cause outages, they can damage sensitive equipment, upset industrial processes, or violate regulatory standards (e.g., IEEE 519 on harmonics). AI offers advanced methods to detect and **classify PQ events** quickly, enabling quicker diagnosis and corrective action.

5.1 Types of PQ Disturbances

- **Voltage Sags and Swells:** Short-duration dips or rises in voltage usually caused by faults, motor starts, or switching operations.²¹
- **Harmonics:** Waveform distortions (frequencies at multiples of the fundamental) often introduced by non-linear loads like variable frequency drives (VFDs) or inverters.²² They can cause extra heating and equipment malfunctions if severe.
- **Transient Overvoltage/Spikes:** Abrupt high-voltage excursions from lightning or switching surges, potentially damaging insulation.
- **Flicker:** Rapid, repeated voltage fluctuations causing visible light flicker, often from arc furnaces or large motor drives.

Identifying the **type and likely source** of PQ events traditionally falls to skilled engineers analyzing waveforms — which can be slow and subjective. **AI classification** can automate this with high accuracy, as many PQ disturbances have distinctive patterns. Machine learning models (like CNNs and Transformer architectures) have demonstrated classification accuracy above 99% on controlled benchmark datasets; field accuracy varies depending on noise levels and whether disturbance types encountered were represented in training data.²³ This means a well-trained model can correctly label events (sag, swell, etc.), enabling quicker responses and targeted follow-up.

5.2 Root Cause Analysis and Action

The value of classification is speeding up **root-cause analysis**. If the AI classifies an event as a "capacitor switching transient", the operator can confirm if a capacitor bank was indeed scheduled to switch at that time (if yes, the event is expected; if not, it might indicate a misoperation to investigate). If a sustained harmonic distortion is flagged and classed as "harmonic likely from customer load", an engineer can direct power quality staff to inspect large commercial customers on that feeder for compliance. Another case: if multiple smart meters report momentary sags concurrently and the AI classifies the pattern as a "feeder fault with auto-reclose", then the utility knows a momentary fault occurred and may decide to dispatch a crew to patrol that feeder segment for likely causes (like a tree branch contact) **even if the breaker did not lock out**. This proactive approach can fix issues before they cause a permanent fault, improving reliability. The AI classification output is integrated into the Distribution Management System (DMS), enabling real-time event labeling within the operator's existing workflow.

5.3 PQ Monitoring Systems

Many utilities already have PQ monitors at key points; AI extends these by not only *detecting events*, but also *identifying them and prioritizing them*. A large facility, for example, might experience dozens of minor sags a month. An AI system can categorize them and perhaps correlate with external factors: e.g., noticing that sags often coincide with the startup of a particular factory's arc furnace at 8AM. The result is a concise report to the engineer — which might prompt working with that customer on mitigation (like adding a soft-start or other techniques). Another example is using ML to detect **waveform distortions indicative of cyber-physical attacks** on the grid (like data injection or rogue inverter behavior); some research has applied generative adversarial networks (GANs) or ensemble methods to detect such anomalies, although they underscore the importance of **interpretability** to trust those results.^{24,25}

5.4 Verification and Regulatory Reporting

Power quality is often regulated or bound by service quality agreements. Thus, any AI classification or measurement likely feeds into compliance reporting (e.g., how many events above certain thresholds occurred). It is critical to **verify the AI's classification accuracy** on historical event logs (where ground truth is known, if possible) and ensure it does not misclassify benign events as serious PQ issues or vice versa. Utility PQ engineers should periodically review a sample of AI-labeled events to verify correctness, akin to calibrating a relay or meter. Over time, the model can be refined with new data (with an auditable procedure for updating it). Document each update to maintain trust (which resonates with **ISO's emphasis on documented information control**).

5.5 Example

A distribution company applied an AI classification model to identify **voltage sags** on feeders and differentiate if they were internally caused (e.g., faults or motor starts) or upstream (transmission supply issues). The model was trained with labeled historical event waveforms and achieved approximately 98% classification accuracy for those two categories. It was integrated into the DMS and, when a sag occurred, it would instantly label it. One day, a significant sag was labeled as "transmission-related". The operator seeing this held off dispatching unnecessary local field crews and instead alerted the transmission operator, saving time and focusing efforts correctly. The sag was indeed traced to a transmission-level capacitor bank mis operation. **This shows AI can expedite root cause identification**, but also the continued role of engineers: the operator needed to trust the classification (backed by prior validation) and make the proper call.

Module 6: Validation Metrics and Reliability Indices — Measuring Success

To responsibly implement AI in grid operations, we need to quantify model performance and track how it improves (or potentially could degrade) operational outcomes. This module focuses on validation metrics for AI models and how their use ties into reliability indices and other KPIs.

6.1 Model Performance Metrics

- **Accuracy and Error Metrics:** For prediction tasks (load, outages, etc.), typical metrics include **MAE (Mean Absolute Error)** and **MAPE (Mean Absolute Percentage Error)** — these tell how far predictions deviate from actual outcomes on average. For classification tasks (e.g., event type detection), metrics like **accuracy, precision, recall, and F1-score** are used. High values (like 97–99% load forecast accuracy or 95%+ in anomaly classification) indicate strong performance, but must be interpreted in context²⁶ — often extreme events have bigger errors, so the engineer must examine performance under stress conditions, not just average cases.
- **Recall versus Precision:** For event detection, **recall** (how many actual events are caught) versus **precision** (how many flagged events are real) is a key trade-off. Depending on the application (safety-critical versus routine), an engineer might tune the model for higher recall (catch all events, risk more false alarms) or higher precision (only flag when fairly certain).
- **Lead Time / Prediction Horizon:** For outage prediction, measure the **average lead time** of correct predictions (e.g., a model might predict outages approximately 42 hours in advance on average²⁷). Longer lead time is valuable but can come at the cost of more uncertainty.
- **AUC (Area Under ROC Curve):** If treating event detection or outage prediction as binary classification, AUC is a good overall measure of discrimination ability. Advanced deep learning models have demonstrated AUC values above 0.95 in published distribution anomaly detection benchmarks, indicating strong discrimination ability.²⁸

6.2 Reliability Indices as Outcome Metrics

Ultimately, if AI in operations is effective, we should see improvements in widely used reliability metrics:

- **SAIDI (System Average Interruption Duration Index):** Measures the average outage duration per customer per year (in minutes). AI that proactively avoids or shortens outages drives SAIDI down. For example, predictive measures in one benchmark scenario reduced SAIDI from approximately 142 to 38 minutes/year²⁹ — utilities should establish their own baselines and track improvements incrementally.
- **SAIFI (System Average Interruption Frequency Index):** Average number of outages per customer. If incipient failures are fixed before causing outages, SAIFI declines. E.g., if a feeder had three expected failures in five years historically, and models catch two of them in advance, those can be turned into planned maintenance (zero-outage events) — thus SAIFI lowers.
- **CAIDI (Customer Average Interruption Duration Index):** The average outage duration for customers who experience an outage. Improving response time through predictive analytics (e.g., staging crews or reconfiguring to isolate faults more quickly) can cut restoration time, thus lowering CAIDI.

From the perspective of **utility regulators and customers**, these indices are key. Thus, any AI solution must be assessed on how it affects SAIDI, SAIFI, etc., not just internal metrics. And importantly, to be credible, any improvements should be traceable to concrete changes (like how many outages prevented by specific actions). **Engineers should gather data on how AI interventions correspond to reliability improvements** — for example, track that "we proactively replaced five transformers flagged by AI in 2025, and those feeders subsequently had no unplanned outages, thereby avoiding an estimated X customer-hours of interruptions and improving SAIDI by Y minutes." These results both justify the AI investment and provide evidence when communicating with regulators or the public.

6.3 Dealing with Uncertainty

No model is perfect. To maintain safety, uncertainty must be accounted for:

- **Prediction intervals and safety margins:** E.g., if a load forecast says 95 MW $\pm$ 5 MW, plan as if 100 MW might occur (especially if that is near a limit).
- **Worst-case scenarios:** For critical decisions (like grid stability under high risk), consider if the AI is wrong. Are fail-safes in place? For example, if an outage risk model fails to predict an outage that then occurs, do we still have emergency crews ready? The answer should be yes — AI augments but does not eliminate the need for emergency preparedness.
- **Continual model evaluation:** Track the model's performance over time. If errors start creeping up (maybe due to changing conditions like new DER patterns), retrain or recalibrate. This is similar to protective relay maintenance — the model is part of the operational assets and needs periodic "tuning" and testing.

6.4 Communicating Uncertainty and Results

When an AI model is used in operations or planning, be transparent about its uncertainty to decision-makers. For example, present not just a predicted number, but also confidence levels or scenario ranges, so that management understands the risk. For instance, "Our outage model forecasts three feeders at high risk for a 1,000+ customer outage this week, with approximately 80% confidence. We will focus crew readiness there, but also prepare backup for other feeders in case of model miss." This realistic communication builds trust because you are **open about limitations and how you are addressing them**.

Module 7: Governance, Model Management and Professional Defensibility

The final module covers the **process and governance** aspects of using AI in grid operations — ensuring each model and its outputs are properly managed, documented, and integrated into enterprise processes in a way that holds up to scrutiny.

7.1 Model Governance and Configuration Control

Just as equipment designs go through approvals and revisions, AI models should too. This involves:

- **Version Control:** Every time a model is trained or updated, assign a version identifier. Document changes (e.g., "Outage Model v1.2 — included new weather data, retrained on 2023 storm data, improved accuracy by 5% on validation"). Use a repository or model management tool. This way, if an unexpected model output contributed to an operational decision, the exact model that produced it is known and can be reviewed.
- **Data Lineage and Audit Trails:** Maintain logs of what data was used to train or feed the model. For example, keep copies of training sets and note sources (AMI data from January 2018–December 2022, NOAA weather data from station X, etc.). Also, log model outputs and any automated decisions triggered. This ensures that later, a forensic analysis can determine how a particular prediction was made and on what basis (an essential step for **professional defensibility** if anything goes wrong).
- **Access and Change Management:** Limit who can modify models and how they are deployed, akin to how only authorized engineers can change protective relay settings. Perhaps require a "change request" process to update a model in production, including testing and management sign-off. This prevents a well-meaning data scientist from inadvertently updating a model in a way that might cause issues in operations.

7.2 FMEA Revisited — Model Use Risks

We already looked at example failure modes (false alarms, missed events, etc.). It is recommended to integrate FMEA or similar risk analysis into the development and deployment of any AI solution. E.g., *What is the impact if the model goes down?* (Mitigation: failover to normal alarms). *What if the model is wrong?* (Mitigation: maintain conservative operations, escalate unusual suggestions to additional review.) This systematic thinking ensures no scenario is overlooked. In practice, combining the FMEA with simulation or pilot testing helps find the threshold for where automation is safe. For instance, some utilities use AI models in a **recommendation mode** to operators first (with the operator required to confirm actions), and only allow automatic actions after a track record of reliability is established and hazard modes ruled out. This incremental approach is a form of risk mitigation.

7.3 Ensuring Regulatory and Ethical Compliance

Utility operations are often subject to regulatory oversight. If an AI is involved in decision-making (say controlling voltage or deciding when to cut off a line for maintenance), there could be regulatory interest in how those decisions are made. Having thorough documentation and rational criteria (and showing a human is supervising) will help satisfy regulators that decisions are being made responsibly. Also, in any **post-incident investigation**, being able to produce logs showing *what the AI recommended, how it was verified by an engineer, and why that was reasonable* is critical. Implementing AI should **raise the bar of diligence, not lower it**.

7.4 Communication and Culture

Internally, foster a culture where engineering teams and operations see AI as a helpful partner, not a mysterious black box. This means **training staff** on how the models work at a conceptual level, what their outputs mean and how to interpret them, and what the limits are. It also means encouraging a feedback loop: if field crews find that an "incipient failure" alert was a false alarm, that feedback should be captured (e.g., marked in the incident system) so the model can learn. Conversely, share success stories — if an AI prevented an outage or diagnosed a power quality issue swiftly, communicate that to increase confidence and adoption.

7.5 Professional Defensibility

Ultimately, if an AI-informed decision is ever challenged (say a major outage occurred that the AI did not predict, or an unnecessary costly action was taken because of a false alert), the responsible engineer should be prepared to defend the actions. By having adhered to a robust process (clearly documented assumptions, thorough testing, risk analysis, etc.), the engineer can show that they exercised due care and that the AI tool was used appropriately as part of an **engineered solution**. They can point to logs, validation results, and the fact that human oversight was exercised at critical junctures. This turns potential suspicion about black-box AI into an evidence-based defense that the approach was systematic and prudent.

In closing, AI techniques like anomaly detection, outage prediction, and power quality analytics are powerful additions to the grid operations toolkit. They enable a **transition from reactive to proactive operations**³⁰ — detecting troubles before customers are impacted. When implemented with careful attention to data quality, model validation, and professional oversight, these tools can significantly **improve reliability indices** and **operational efficiency**, all while respecting the paramount charge of engineers to ensure **public safety and grid stability**. With the knowledge from this course, participants can begin to responsibly integrate AI into grid operations, enhancing their ability to maintain a resilient and high-quality power supply in an increasingly complex energy landscape.

References

1. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
2. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
3. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
4. Ernst & Young LLP. "AI Can Help Utilities Predict Grid Outages." https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages
5. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
6. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
7. Ernst & Young LLP. "AI Can Help Utilities Predict Grid Outages." https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages
8. Ernst & Young LLP. "AI Can Help Utilities Predict Grid Outages." https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages
9. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
10. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
11. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
12. Frontiers in Energy Research. "Anomaly Detection in Power Distribution Systems." <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1364456/full>
13. Frontiers in Energy Research. "Anomaly Detection in Power Distribution Systems." <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1364456/full>
14. MDPI Applied Sciences. "Power Quality Disturbance Classification Using Deep Learning." <https://www.mdpi.com/2076-3417/15/3/1442>
15. Ernst & Young LLP. "AI Can Help Utilities Predict Grid Outages." https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages
16. Ernst & Young LLP. "AI Can Help Utilities Predict Grid Outages." https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages
17. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
18. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
19. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>

20. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
21. MDPI Applied Sciences. "Power Quality Disturbance Classification Using Deep Learning." <https://www.mdpi.com/2076-3417/15/3/1442>
22. MDPI Applied Sciences. "Power Quality Disturbance Classification Using Deep Learning." <https://www.mdpi.com/2076-3417/15/3/1442>
23. MDPI Applied Sciences. "Power Quality Disturbance Classification Using Deep Learning." <https://www.mdpi.com/2076-3417/15/3/1442>
24. Frontiers in Energy Research. "Anomaly Detection in Power Distribution Systems." <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1364456/full>
25. Frontiers in Energy Research. "Anomaly Detection in Power Distribution Systems." <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1364456/full>
26. MDPI Applied Sciences. "Power Quality Disturbance Classification Using Deep Learning." <https://www.mdpi.com/2076-3417/15/3/1442>
27. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
28. Frontiers in Energy Research. "Anomaly Detection in Power Distribution Systems." <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1364456/full>
29. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>
30. Energy Solutions. "AI in Grid Management: Predicting Blackouts." <https://energy-solutions.co/articles/sub/ai-grid-management-predicting-blackouts>