The Promise of Automated Inspection in Electric Utilities

White paper | systemswithintelligence.com | info@systemswithintelligence.com

Time-based inspection built the reliability record of the modern grid, and it is still the basis for many utilities. But, on its own, the time-based approach is no longer sufficient for many substations.

The strategic question for any utility is how to design a program that matches the right technology to the right inspection task, and then how to build that program for the future.

Looking ahead, we cannot help but be excited about the promise of automated inspections and the revolution in continuous inspection that AI will help to enable.

Why Automated Inspection?

CIGRE Technical Brochure 642, the international transformer reliability survey that analyzed 964 major failures across 21 countries, found that chronological age is a weaker predictor of failure than most replacement strategies assume.

Transformers in slightly degraded condition fail under unusual operational stress. This means technical condition is more important than calendar age, and inspection frequency matters more than consistency with a published schedule.

Yet increased frequency can only be part of the solution. Manual inspection is inherently variable. Inspection findings depend on the inspector's experience, load and lighting conditions, equipment access angles, and even the discipline with which documentation is handled. In other words, even two inspectors examining the same asset on the same day may record different findings, rendering any trend analysis unreliable. This is perhaps why estimates show that even two best-in-class inspectors can catch only about 96% of defects.

Scale is another issue. The ratio of assets to qualified inspection personnel is worsening as the industry's experienced workforce retires, with the US DoE expecting 40% of the power industry to be retirement-eligible by the end of the decade.

At the same time, electrification is accelerating. Commercial data centers compete with residential EV chargers for grid capacity that is already stretched thin, with data centers expected to account for 55% of peak load growth, or 90GW, over the next five years, and EV charging projected to represent as much as 4.6% of US power demand by 2030. In other words, there is more to inspect with fewer people to inspect it, and higher stakes if something slips through.

These problems converge in one place: the Tier 2 and Tier 3 substations built forty, fifty, or sixty years ago, which simply were not designed for the loads today's grid puts on them and which face increasing strain as they age out. (The problem is so bad that the American Society of Civil Engineers (ASCE) gave U.S. energy infrastructure a C-minus grade in a recent report.)

For any one of these reasons, let alone for all four, a paradigm shift in inspection is desperately needed.

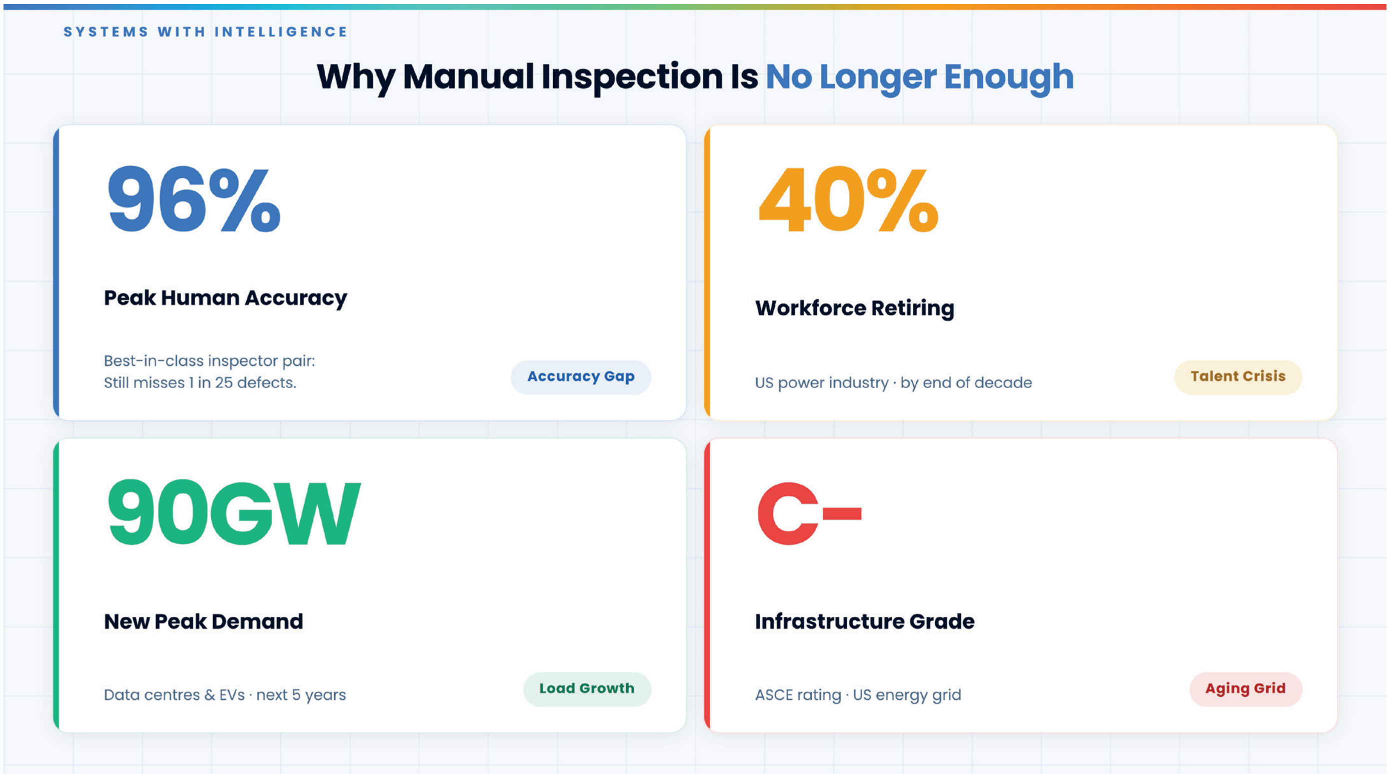

But First... What "Is" Automated Inspection?

When defining what we mean by "automated inspection," it's important to discuss the difference between two fundamentally different kinds of data, only one of which is really the purview of AI.

The first is *deterministic* data, which is precise and auditable. Equipment can monitor deterministic data and alert operators if thresholds are reached. No probability is involved.

This is the approach codified in IEEE Std 62 and IEEE C57.152, the standards that have anchored thermographic field testing for decades. Although artificial intelligence can help with thermal inspection (as we will discuss below), it does not require an AI layer to be useful.

By contrast, automated visual inspection relies on heavily on *probabilistic* data, which is not so much measured as predicted. For tasks like reading gauges, detecting leaks, and identifying changes in equipment condition, an AI machine learning model can be trained to recognize what "healthy" looks like and flag deviations from that benchmark.

This essentially makes the outputs of a probabilistic system guesses, even when they appear certain. When AI reports "this gauge appears to be reading too high," it is really saying something like "I am 90% sure this reading is different from before, and the difference likely means there is a problem."

Neither deterministic nor probabilistic data is "better." Both kinds of information belong in an automated inspection program.

But when we speak of "automated inspection," the difference between deterministic thermal data and probabilistic visual data means the phrase is usually shorthand for "visual inspection enhanced and automated by AI."

That is different from remote video monitoring, which only moves the human eye to a screen in a control room. Neither is alarm-based SCADA, which catches conditions that have already crossed a defined limit.

Put simply, automated visual inspection is the scheduled, systematic collection of data at defined inspection points, combined with automated analysis by artificial intelligence that produces structured findings an operator can act on.

And although the technology is currently less mature or validated than thermal inspection, AI is developing quickly, and to good effect. With as little hyperbole as possible, automated visual inspection is very likely a key to the future of the industry.

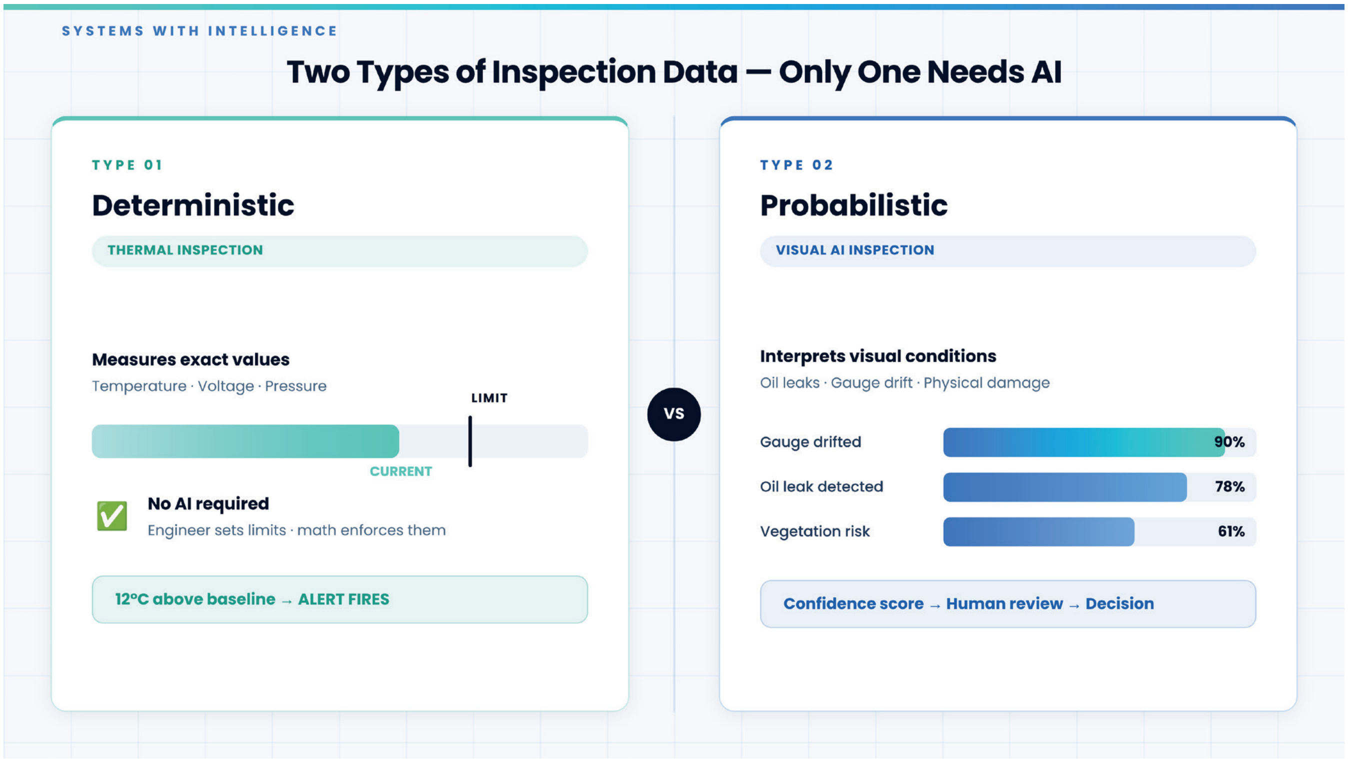

Getting Visual Data in the First Place: Fixed Cameras, Drones, and Robots

Three delivery technologies dominate the conversation about automated visual inspection: fixed installed cameras, unmanned aerial drones, and ground-based robots.

Cameras

For continuous inspection of known critical assets, the most mature option is fixed or pan-tilt-zoom cameras, permanently mounted at strategic positions to cover defined inspection points. They are always on, and unlike other technologies, they don't require a deployment window, a flight authorization, or a weather check.

More advantageous, they return to the identical camera position and zoom level every inspection cycle. (This matters more than it might seem at first, as we explain below.)

In an automated scenario, cameras cycle through preset positions, capturing images at a defined cadence, and then feed those images to a reporting environment that organizes and classifies them. An operator opens an asset-centric interface that presents historical images with anomalies flagged for review. Inspection is continuous, and documentation is automatic, and the operator's role shifts from performing inspections to reviewing findings. (SWI's INTELLINSPECT™ platform is one operational implementation of this kind of approach.)

Realistically, cameras do suffer from obvious limitations in coverage. They can only inspect what they can see from their installed positions; accordingly, blind spots are a design challenge that requires careful camera placement. Even with well-chosen installations and industry-leading PTZ capability, a fixed camera can't get behind bus work, read a gauge from the wrong angle, or see the underside of equipment.

For this reason, practical fixed-camera programs typically target 85 to 90 percent coverage of critical assets, with the simple "given" that some inspection tasks belong elsewhere.

Drones

Aerial drones provide flexibility for inspecting transmission lines, large substation structures, and post-event damage. Fixed cameras are always mounted above the asset so they can also see the top of the asset. This is not a particular advantage of drones. They can inspect the underside of equipment, and survey large areas in ways no fixed camera can match.

However, their constraints are just as real. Drones require licensed pilots and typically require authorization to fly; as a consequence, they remain heavily regulated in most jurisdictions.

They are also vulnerable to the electromagnetic interference that substations generate in abundance, which makes GPS and wireless communication less reliable near high-voltage equipment, and they are weather-dependent in ways fixed cameras are not.

Additionally, because each flight captures data at slightly different angles, altitudes, and distances, drone-collected data poses a serious challenge for AI-based trend analysis.

Finally, the per-substation case for a dedicated drone is also hard to make. Most practical drone programs work on a fleet basis, with a crew and equipment deployed across many sites. That constraint alone makes drones best positioned as a complement to continuous visual monitoring rather than a replacement for it.

Robots

Robots for inspection look promising, although their deployment in substation environments remains limited compared to cameras and drones.

Their strongest advantage is integration: a purpose-built utility robot can carry thermal, visual, acoustic, and LiDAR sensors, and even perform gas detection. A programmed path can deliver inspection consistency similar to that of cameras, though rough terrain (and snow and ice), introduces uncertainties that neither drones nor cameras face.

Robots are also in the early stages of development, and their usefulness is constrained by battery life, weather sensitivity, and unit cost. They may fit a large site where ground-level data collection justifies the investment, but "set it and forget it" robot inspection is a long way away.

CRITERION	FIXED INSTALLED CAMERAS	DRONES	GROUND ROBOTS
Inspection frequency	Continuous or on schedule, no deployment needed	Constrained by pilots, weather windows, and airspace approvals	Constrained by battery cycles and terrain
Positional consistency	Same position, same zoom, every cycle	Changes slightly with every flight, which makes AI trend analysis hard	Good path repeatability. Terrain can introduce sensor-position shifts
24/7 operation	Yes, in all weather	No. Visibility, wind, and airspace rules all interfere	Limited; depends on battery and weather
AI anomaly detection	Works well; consistent angles make baseline comparison reliable	Still maturing. The flight-to-flight variability is the main technical obstacle	Still maturing. Multi-sensor payloads add capability but complexity
Access to asset detail	Only what the installed positions can see; blind spots need planning	Can reach almost any angle. Fixed cameras can also see tops of assets	Ground-level only. No vertical reach
EMI hardening	Built for permanent installation in high-EMI environments	Vulnerable. EMI disrupts GPS and wireless communication	Mixed. Purpose-built utility robots are hardened; off-the-shelf ones are usually not
Regulatory requirements	None after installation	Licensed pilots, flight plans, and airspace authorization	Site access protocols, but nothing from regulators
SCADA/APM integration	Hardwired to site network	Happens after the flight. Data has to be uploaded and processed	Usually wireless. Latency and reliability vary
Capital cost	Higher upfront, low to operate	Lower hardware cost, but pilots and operations add up	High. Purpose-built utility robots are expensive
Deployment timeline	Longer time required to install and commission	Fast. No installation required	Moderate. Needs a site survey, path programming, installation of robot 'garage'
Best suited for	Continuous monitoring of critical substation assets	Inspection of towers and lines	Large sites where ground-level patrols are worth the investment

As the chart above demonstrates, every tool has its application. Each of the three approaches addresses the same underlying challenge of collecting image data without sending a person to a site, yet none can solve the problem on its own.

That is why we believe that a mature automated inspection program combines approaches: fixed cameras for continuous monitoring, drones extending coverage where fixed cameras cannot reach, and robots adding ground-level capability where site scale justifies it.

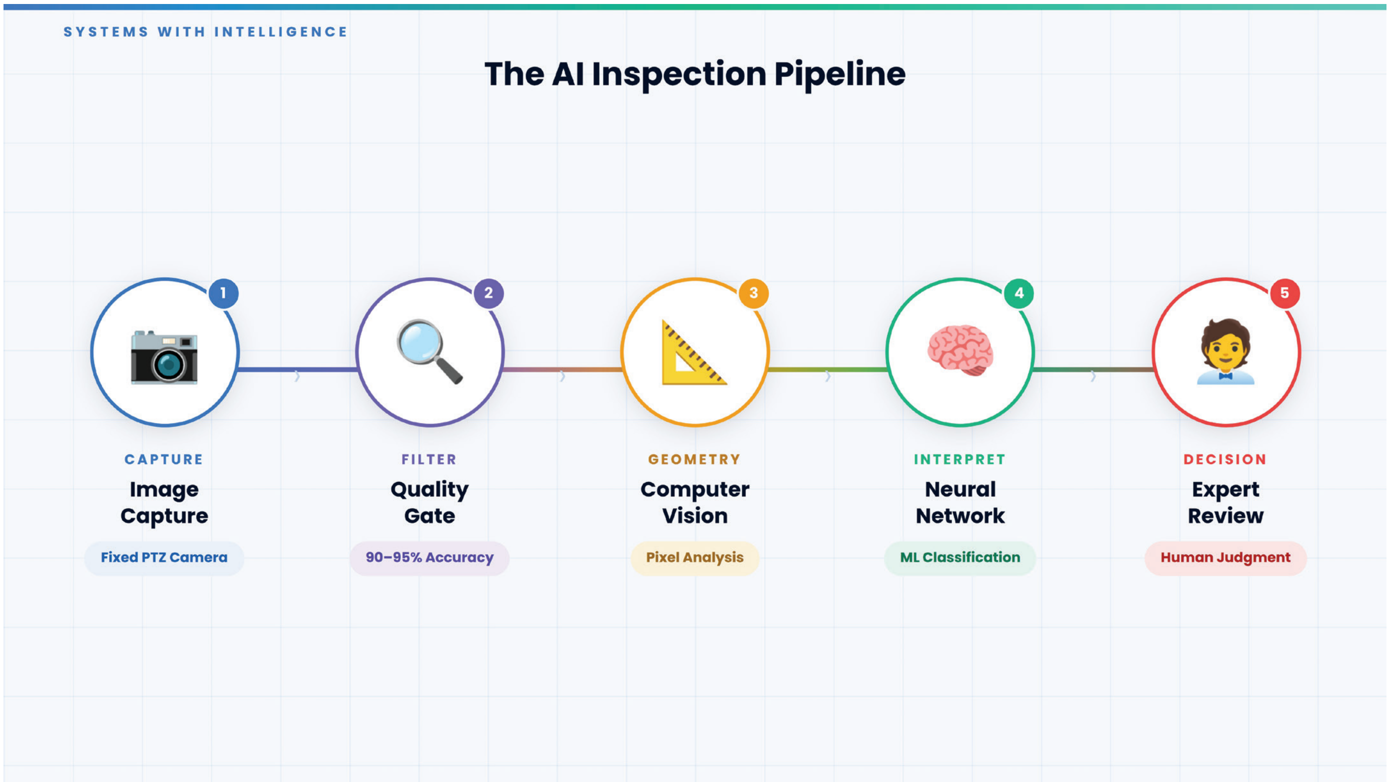

Automating Visual Data Analysis with AI

Artificial intelligence in the utility space usually involves one or more of three different technologies, or "models."

The first is what is most often meant when people talk about machine learning: a classic neural network in which labeled data is passed through a training loop until the AI learns which features correlate with which outputs. The result is a pattern recognizer.

For tightly scoped problems, this pattern-recognition approach can produce good results with surprisingly little data. (For example, a confined sub-problem, such as "Is vegetation of any kind, no matter what it looks like, crossing this line?" can work with as few as 200 training images.)

The second is so-called "computer vision." This is closer to image processing, where the system does pixel-by-pixel interpretation against known geometric features. "Compare this frame to the previous frame and flag the pixels that changed." Interestingly, computer vision pipelines are deterministic in a way neural networks are not: someone sets the thresholds, the math does the work, and the output is predictable.

The third is a hybrid model that combines both approaches.

Consider gauge reading as a clear example. Reading a gauge sounds simple until you realize that dials come in every configuration, needles come in every thickness, gauge labels are far from standard, and the glass on a gauge could be dirty or obscured by reflections. It quickly becomes clear that a purely computer-vision approach to gauge reading would collapse under the weight of all the variations.

Yet a neural-network approach would struggle with the geometric precision required to read a value.

The solution? Let computer vision handle the geometric analysis and let a machine learning model handle the interpretation and candidate selection. Breaking the hard problem into smaller, well-defined sub-problems makes each piece easier to tune and validate.

The approach is data-intensive, however. A generalized task like gauge reading across many gauge types might start with 3,000 to 5,000 raw images and expand to 15,000 or 30,000 through augmentation: shifting, skewing, and varying the lighting and contrast of existing images. (At Systems With Intelligence, roughly 30 percent of our training data is synthetic, used for augmentation and benchmarking against scenarios that are hard to capture live. The rest comes from real installations at customer sites.)

And to be clear, this specialized training is absolutely necessary, both because the stakes are high and because consumer AI is simply not up to the task. Our own internal prototype testing showed accurate readings in around 30 seconds, where the most popular consumer AI available today took upwards of 15 minutes to confidently produce an entirely wrong answer.

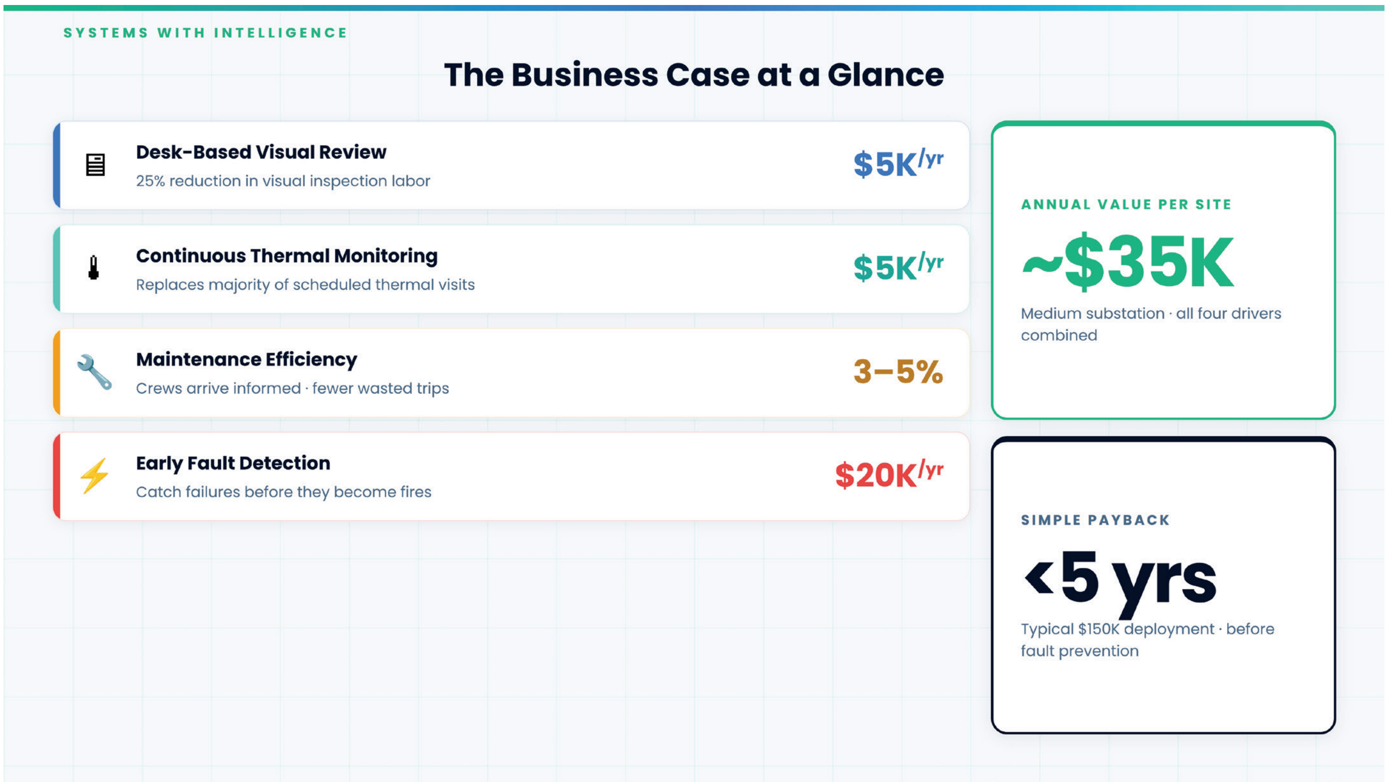

Where AI Adds Value (and Where it Does Not)

Understandably, AI adds more value to visual inspection than to thermal imaging, because thermal data is already deterministic.

A thermal camera provides temperature values at the pixel level, and basic comparative math captures most of what matters: phase-to-phase differentials and trending against historical readings, all without probability or confidence scores.

Where AI will eventually add value to thermal is in looking beyond the defined regions of interest that most analytics rely on today, scanning the full thermal image for anomalies that an operator did not think to measure.

Artificial intelligence also opens the door to cross-validating thermal information against visual data from a companion camera. A thermal signature without a visible cause is more interesting than one that correlates with a known operational event, and an AI that can make that correlation is more useful than either sensor alone.

However, it's predominantly on the visual side that AI earns its place. Many of the failure precursors that matter in substation equipment are visual: oil leakage, insulator contamination, gauge drift, animal intrusion, vegetation encroachment, physical damage after weather events, and so on.

These are the tasks where pattern recognition against a consistent baseline produces real operational value, and where AI stands to make great strides.

Data Consistency is the Prerequisite

Regardless of which AI model is used to return what data, the success of visual AI in a substation depends entirely on the consistency of the input data.

A model trained to identify a thermal hot spot on a bushing needs to see that bushing from the same angle and zoom level, under comparable lighting and load conditions, in every inspection cycle. Every dimension of variability the model has to accommodate is a dimension along which it can be wrong.

Lighting, weather and load conditions matter for the same reason. That same bushing at 40 percent and 90 percent load produces very different thermal signatures; analytics that do not account for this will either miss real problems or mistakenly flag normal operation.

A more reliable approach would train models with load and environmental context built in, so the model learns that a given reading is acceptable in winter at noon with moderate load and concerning in summer at midnight with light load. This is where the field is moving, but success with AI here is still a frustrating dream.

Fault Classification: A Particular Problem

The most ambitious version of AI inspection is a system that detects an anomaly, classifies the fault, and recommends a course of action. This is often called the Holy Grail of the category because, in this respect, AI is years away from being appropriate for deployment in critical infrastructure.

This is partly for technical reasons. Any probabilistic system produces both false positives and false negatives. False positives erode operator trust by crying wolf, while false negatives defeat the purpose of the inspection by missing real problems, but the trade-off is unavoidable. Lower the detection threshold, and you catch more real issues at the cost of more false alarms. Raise it, and you get the reverse. There is no threshold setting that eliminates both.

Reducing errors

Mature implementations reduce errors through structural engineering choices rather than through threshold tuning alone. Three techniques are worth discussing.

The first is a two-stage pipeline with a binary classification gate. The idea is to ask a simple yes-or-no question before attempting the harder task: does this image contain the thing the model is looking for at all? A binary gate can be trained to 90 or 95 percent accuracy with modest data, and it prevents the classification stage from producing confident outputs on low-quality images. (A common failure mode in naive single-stage pipelines, and one that is largely solvable.)

The second is persistence thresholding. Rather than reporting any anomaly the moment it is detected, this kind of system requires that an anomaly appear consistently across multiple inspection cycles before it surfaces as a finding. This filters out transient noise: birds crossing the frame, sun reflections on equipment, momentary lighting changes, brief weather artifacts, and so on. Thresholding trades a relatively small amount of detection latency for a large reduction in the volume of false alarms, but this concession is appropriate when most of the degradation that matters in a substation develops over days, not seconds.

The third approach, as we have hinted above, involves breaking issues into more defined sub-problems.

For example, think again of our gauge-reading challenge. A gauge-reading pipeline that first locates the gauge, then finds the needle, then reads the needle value relative to the label geometry, is far more reliable than one that produces a reading in a single step. In this way, breaking a hard problem into smaller, well-defined sub-problems becomes itself a false-positive reduction strategy, because each sub-problem can be tuned and validated independently.

Taking accountability

Liability is still the much larger challenge with probabilistic systems. When an asset expert identifies a fault, he or she is typically held accountable for that judgment. When AI determines there is a fault, on the other hand, accountability becomes ambiguous. Was a missed failure the fault of the operator or of the AI? Who would share responsibility for an hours-long outage? The utility? The AI developer? The consultant who put the system in place? The team that trained it?

We believe the right role for AI is to surface anomalies to the asset experts who are accountable for the decisions, because fault determination is expert work. The job of the inspection system is to make that expert work faster and better-informed, not to replace it.

Building the Business Case for AI

Every utility has stories about catches that came in time. Some have (horror) stories about ones that didn't. Both kinds bookend the business case for automated inspection.

Start with the money. An inspector reviewing current and historical imagery from a desk, without a truck roll, cuts annual visual inspection labor by around 25 percent. Continuous thermal monitoring does more, taking over most of the thermography work on the assets it covers and saving another 25%.

Plus, pre-work visibility into site conditions means crews show up loaded for the job in front of them, which is worth three to five percent on planned and corrective maintenance.

But the bigger number is harder to predict. Catching a developing transformer fault before it becomes a fire is more than enough to justify investment in continuous inspection.

Most sites will not see a catch every year, but across a fleet, even a two-percent annual find rate averages to roughly $20,000 per site per year.

On a medium-sized substation, these four drivers together add up to about $35,000 a year in monetized savings. Against a typical $150,000 deployment, the simple payback comes in under five years. That is the spreadsheet answer, but the spreadsheet misses the rest of it.

Operators who start reviewing continuous inspection data all say the same thing: they find issues they simply did not know existed. A slow thermal drift, a gauge reading sliding the wrong way, a vibration pattern nobody had a reason to look for... None of these factors into an ROI calculation. But these are precisely the kinds of findings that change how a utility thinks about its fleet.

They're the reason automated inspection tends to stay once it arrives.

SYSTEMS WITH INTELLIGENCE

Three Delivery Technologies — No Single Answer

OPTION 01

Fixed Cameras

Continuous Monitoring

01

OPTION 02

Drones

Periodic Inspection

02

OPTION 03

Ground Robots

Large-Site Patrol

03

	Cameras	Drones	Robots
Always On	✔	✘	~
Aerial Access	✘	✔	✘
AI Trend Analysis	✔	~	~
Multi-Sensor	~	~	✔
No Operator Required	✔	✘	~

LEGEND | ✔ Yes — fully supported ~ Partial — limited or emerging capability ✘ No — not applicable

The Practical Question

For a utility approaching automated inspection for the first time, the highest-value starting point is usually in Tier 2 and Tier 3 substations. They are where the inspection gap is widest, where the risk is growing fastest, and where the ROI is cleanest to demonstrate.

From there, you will need to match the right delivery technology to each inspection task, with fixed cameras handling continuous monitoring of critical assets, drones reaching what fixed cameras cannot, and robots covering large sites where ground-level data collection justifies the investment.

Unfortunately, the AI hype cycle means you will also face issues of vendor enthusiasm. A partner that acknowledges what the technology can and cannot do today is more valuable than one promising an “easy button” that does not exist. There are no shortcuts and no perfect fits; any vendor claiming otherwise should be evaluated skeptically.

But regardless of these challenges, we believe strongly that exploring AI automation makes sense. Automated inspection is not about replacing what utilities are already doing well. It is about filling a gap that manual inspection alone can no longer address, using technology that solves problems instead of creating more.

As Angelo Rizzo, Systems With Intelligence president and CEO, puts it: “Utilities are under pressure to do more inspections and review more data with fewer resources. Automated inspection helps operators remotely assess the condition of assets with greater accuracy and efficiency, supporting a smarter approach to grid maintenance.”

None of this is easy, and none of it is finished. But emerging tools are available today, and the use cases are defensible. And crucially, the utilities that start asking the questions now will be the ones best positioned for what comes next.

Systems With Intelligence has been enabling automated inspection systems at electric utilities since 2009, with over 900 customers worldwide. To discuss automated inspection program design for your utility, contact sales@systemswithintelligence.com or visit www.systemswithintelligence.com.

About Systems With Intelligence

Systems With Intelligence (SWI) is the trusted global leader in Touchless™ Monitoring solutions for electric utilities and industrial applications. Founded over 15 years ago, SWI pioneered real-time thermal imaging for substations and has deployed thousands of monitoring systems worldwide. The company's sensor networks and analytics platforms enable utilities to achieve dramatic O&M savings while building the data foundation for grid modernization and autonomous operations.