Computer Vision - AI for Construction

Course No: C01-039
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-First Mindset & AI's Role

Overview: Construction sites are inherently high-risk environments that demand rigorous QA/QC and safety oversight. PEs working in construction have a duty to ensure structures are built correctly and safely, as lives depend on it.

1.1 Construction QA/QC and Safety – Why We Need Better Tools

Construction remains one of the most hazardous industries. It accounts for about 20% of workplace fatalities in the U.S. despite only ~7% of the workforce.[1]

Traditional QA/QC and safety practices rely on periodic manual inspections, checklists, and post-incident learnings. These methods are effective up to a point but can be labor-intensive, subjective, and not continuous. For example, a rebar inspection often involves an engineer manually measuring spacing and comparing to drawings – a meticulous, time-consuming process prone to human error if rushed.

Safety professionals conduct site walks and rely on workers to report near misses – but some near misses go unnoticed or unreported. In one analysis, near-miss reports increased by 75% after deploying an AI safety system, indicating many more hazards were caught than previously.[2]

AI-driven computer vision can shift these processes from reactive to proactive. By monitoring jobsites 24/7 via cameras and analyzing footage in real time, AI can spot issues as – or even before – they occur. For instance:

- Noticing a misplaced rebar or missing anchor bolt on a beam before concrete is poured, prompting immediate correction (instead of a costly post-pour fix).
- Continuously checking if workers are wearing required PPE like hard hats and safety vests, sending alerts if not.
- Monitoring a busy equipment zone and alerting when a worker and a vehicle come dangerously close (a potential near-miss collision) so preventive action can be taken.

A study focusing on AI safety in construction projects reported a 34% reduction in total safety incidents over one year after implementing CV-based PPE and hazard monitoring.[3]

Likewise, QA issues caught preemptively mean fewer structural deficiencies or schedule delays, which translates to safer, more reliable outcomes.

1.2 AI as Decision-Support, Not Decision-Maker

Despite these benefits, it cannot be overstated: The engineer is still in charge. PEs integrating AI vision tools must ensure they augment human judgment, not replace it. According to engineering ethics (e.g., NSPE and ASCE guidelines), the licensed engineer must hold paramount the safety of the public and cannot delegate this responsibility to an algorithm. In practice, this means:

- If a computer vision system flags a QA issue or safety hazard, an engineer or qualified professional should verify the issue and decide on the action (like an inspection or work stoppage).
- Engineers should treat AI outputs as one source of evidence, alongside drawings, specs, physical observations, and tests, when making decisions.
- The PE must document how AI outputs were used and ensure any final certification or report is under their direct control and review (the AI does not "sign off" on anything).

By framing AI as a smart assistant – a constant watcher that never gets tired – PEs can extend their oversight capabilities. It is similar to how a structural engineer uses a finite element program: the software performs intensive numerical calculations, but the engineer interprets the results and ensures they are reasonable. Here too, AI does heavy monitoring while the engineer interprets the findings, confirming they align with reality and standards.

1.3 Multidisciplinary Relevance

Implementing computer vision in construction touches multiple engineering disciplines:

- **Civil/Structural Engineers:** Oversee building and infrastructure projects; they can use AI to automatically check construction quality (e.g., ensuring rebar spacing meets design), to track schedule progress, and to enhance site safety observation.
- **Mechanical Engineers:** Heavy machinery (cranes, excavators, HVAC equipment) is central on jobsites. Mechanical engineers can benefit from AI that monitors equipment operation – e.g., a vision system detecting when a crane's swing path is clear or when maintenance is needed based on visual cues like smoke or leaks.
- **Electrical Engineers:** Support construction tech via design of camera and network systems. They may also use vision tech for electrical QA (verifying correct cable routing, or thermal cameras for loose connections) and ensure devices on site meet electrical safety requirements.
- **Can leverage AI vision for environmental compliance on construction sites, such as monitoring dust generation, identifying erosion control measures (e.g., silt fences in place), or detecting hazardous material spills.** Environmental Engineers:

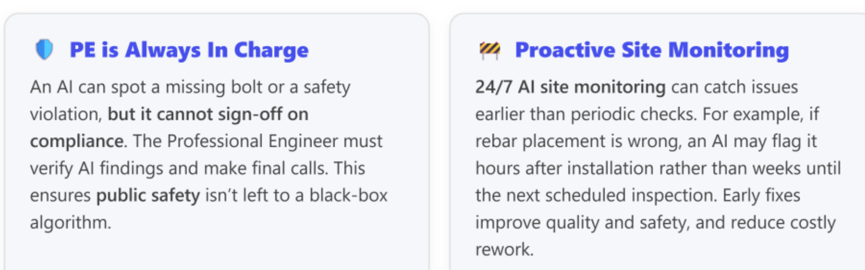

Knowledge Check

1. Why might an AI model trained on day-time site images perform poorly at night?
2. Give one example of occlusion on a construction site that could trick a computer vision system.
3. How can overlapping camera coverage help mitigate some vision system limitations?

Module 2: Computer Vision Fundamentals & Model Limitations

Overview: Computer vision (CV) is the field that enables machines to interpret image and video content. In construction, typical CV tasks include object detection and activity recognition. This module provides a primer on how CV works and, importantly, discusses limitations and challenges.

2.1 How Computer Vision "Sees" a Construction Site

Modern CV systems often use deep learning, specifically convolutional neural networks (CNNs), trained on many labeled images to recognize patterns. Some typical techniques in this domain:

- **They output bounding boxes around specific items of interest (hard hats, safety vests, machinery, etc.). These models scan each frame of video and classify parts of it into known object categories.** Object Detection Models (e.g., YOLO, Faster R-CNN):
- **Used if we need to highlight regions (e.g., distinguishing floor vs wall or identifying regions of concrete vs rebar). It assigns a class label to each pixel in an image, which could help in identifying extents of a concrete pour or area of installed rebar.** Semantic Segmentation:
- **Some safety applications track human poses (skeleton estimation) to detect risky movements, like someone bending over a dangerous location or climbing a ladder without fall protection.** Pose Estimation & Activity Recognition:

Photogrammetry & 3D Vision: For progress monitoring, images from different angles or drones can be stitched into 3D models or compared with BIM data. Advanced systems can compare site reality to the 4D BIM schedule to deduce progress; one AI platform achieved ~98% accuracy in matching installed structural elements to the BIM model *(per vendor data)*.[6]

Computer vision can also automatically count installed components like rebar rods or concrete panels, providing real-time quantity updates (saving days of manual surveying).[7]

2.2 Common Computer Vision Challenges on Construction Sites

Lighting Variations: Sites can be brightly lit under the sun or dim at dawn/dusk. Many CV models trained on normal daylight images see a massive drop in accuracy in low-light conditions – studies have shown that CV models trained primarily on daylight images can suffer dramatic accuracy losses in low-light or nighttime conditions — mitigation through IR-capable cameras and low-light training data is essential.[8]

- Mitigation: Provide adequate lighting or use cameras with night-vision (IR) and ensure the AI is trained on low-light data.

Occlusion (Obstructions): CV models can miss heavily occluded objects – in tests, detection error rates have been observed to increase dramatically (5× or more accuracy degradation) under heavy occlusion compared to clear views.[9]

Mitigation: Use multiple camera angles (overlapping coverage) so if one angle is blocked, another sees the target – overlapping camera placement is recommended for critical zones.[10]

Background Clutter: Construction sites have lots of moving parts. A vision model that is not robust might confuse one thing for another in a busy scene. Training with real construction imagery and applying detection filters can reduce such errors.

Camera Perspective & Resolution: A camera placed far away might not catch fine details. Distortion from angle (looking at something obliquely) can also affect detection. Mitigation: Place cameras thoughtfully – use zoom or high-res cameras for detail-oriented checks, and mix static wide-angle cameras with mobile cameras (drones or wearables) for close inspections.

Weather and Site Conditions: Rain, dust, or snow can obscure camera lenses or objects. Ruggedizing cameras (wipers, cleaning schedule) and using weather filters or sensor fusion can help. If cameras are down or obscured, ensure critical tasks have a backup human check.

2.3 Model Bias & Training Data Limitations

What the model "knows" is only what it has seen in training. If all training images show brand-new shiny rebar but the site has rusted or epoxy-coated rebar, the model may initially struggle to recognize it. Similarly, a PPE detection model trained mostly on images in daylight with yellow hard hats and green vests might be less accurate for a worker wearing an orange vest at dusk with dust on it. Retraining or fine-tuning models on actual site imagery improves performance.

Knowledge Check

4. Why might an AI model trained on daylight site images perform poorly at night, and what is one mitigation?
5. Why might an AI model trained on daylight site images perform poorly at night, and what is one mitigation?
6. How might AI-generated alerts be integrated into a project's existing quality or safety management tools?

Module 3: False Positives, False Negatives, and Performance Trade-offs

Overview: No AI system is perfect – they can raise alarms when nothing is wrong (false positives) or miss real issues (false negatives). Balancing these errors is crucial for construction QA/QC and safety.

3.1 Definitions and Significance

False Positive (FP): The system indicates a problem when there is not one. Example QA FP: AI thinks a rebar is missing but it is actually there (perhaps covered by a shadow). Example safety FP: AI thinks a worker is not wearing a glove because it cannot see the hand clearly. FPs can cause unneeded interruptions or alarm fatigue if frequent.

False Negative (FN): The system fails to detect an actual problem. Example QA FN: AI misses a missing anchor bolt and concrete is poured – potentially a serious structural issue left unaddressed. Example safety FN: AI overlooks a worker not wearing a harness who later falls. FNs are more dangerous in safety scenarios because they translate to accidents.

In general, safety monitoring systems prioritize minimizing false negatives even if it means some extra false positives, because missing a real hazard could lead to injury. The optimal threshold needs adjusting through pilot tests.

3.2 Performance Metrics (Precision, Recall) in this Context

Precision: Out of all issues the AI flagged, how many were actual issues? Low precision (many false alarms) reduces trust in the system.

Recall (Sensitivity): Out of all actual issues, how many did the AI catch? We want high recall to not miss critical safety events. That might mean accepting some false positives.

Accuracy: The overall fraction of correct outputs. Accuracy can be misleading if normal conditions dominate; precision/recall are more useful metrics.

False Alarm Rate vs Miss Rate: Some practitioners prefer terms like "false alarms per hour" to gauge manageability.

3.3 Tuning and Trade-offs

Usually, the sensitivity of detection can be tuned by adjusting a confidence threshold in the AI model. If the threshold is low (model is very eager to alert), it flags anything remotely suspicious (high recall, low precision). If set high, it flags only very certain issues (high precision, lower recall).

Engineers, in collaboration with site managers and safety officers, should decide where to set these thresholds. A typical approach is to start with conservative thresholds during a pilot phase, then gradually tighten to reduce false positives as confidence in the system grows – consistent with ISO 9001 continuous improvement.

3.4 False Alarm and Missed Detection Examples

Safety (PPE): Early experiments with PPE detection have reported high accuracy (90%+) but occasional errors in trickier conditions. Requiring the system to see a person in multiple consecutive frames without a helmet before alerting reduces one-off false positives.

Quality (Rebar): If an AI analyzing a rebar mat sees a dark spot (shadow or stain) and thinks it is a missing rebar, that is a false positive. Mitigation: Adjust the camera angle or add lighting to reduce confusing shadows.

Safety Near-miss (Missed Detection): At a highway construction site at night, a dump truck reverses dangerously close to a worker but poor lighting means the AI does not register it. Mitigation: Add more cameras or sensors (proximity sensor on vehicles), improve lighting, or position cameras closer to critical zones.

Quality Pour Readiness (Missed Detection): The AI fails to detect that a small embedded conduit was not placed before a slab pour because it was not visible from the camera angle. Mitigation: Never rely solely on one camera angle for critical checks; always include a human verification sweep.

Knowledge Check

7. What is worse in a safety monitoring context: a false positive or a false negative? Why?
8. If an AI model has very high recall but low precision for detecting hazards, what does that mean in practice?
9. How might you reduce false positives in a rebar detection system if shadows are causing confusion?

Module 4: Evidence Capture and Audit Trail

Overview: One big advantage of AI-driven vision systems is automatic evidence collection. Every time the system analyzes a scene it can save relevant snapshots, annotated images, or logs. This module discusses how such evidence can be used for QA/QC documentation and safety audits.

4.1 What is Captured and Why

When a CV/AI system monitors a site, it can be configured to save evidence for certain events or on a schedule:

- Quality Milestones: When the system verifies rebar placement in a slab before a pour, it can store images or annotated screenshots showing the rebar grid with a timestamp and overlay confirming requirements met. This becomes part of the project QA records.
- Issues and Alerts: For any anomaly the AI flags (e.g., missing component, safety violation, near-miss), the system can save an image or short video of that moment. This is critical for later verification and can go into a Non-Conformance Report (NCR) or incident report.
- Continuous Logs: Some systems keep periodic snapshots (e.g., hourly site photos) as general documentation of site progress and conditions.

4.2 Traceability and QA Documentation

Integrating captured evidence into your Quality Management System (QMS) ensures traceability. ISO 9001 emphasizes maintaining records of inspections and ensuring changes are traceable. With AI:

- Each automated check produces a digital record (image and metadata like time and location) that can be stored with other QA documentation – akin to a signed inspection form.
- If an auditor asks, "How do you know the rebar was placed correctly before pouring the concrete in section X?", you can retrieve the AI-inspected image from that day showing the layout, plus any notes that the engineer resolved flagged issues.
- On the safety side, if regulators inquire about proactive safety measures, you can show records of AI-detected near-misses and the corrective actions taken with video evidence.

4.3 Using AI Evidence in Investigations

If a safety incident or near-miss does occur, the AI's archived video can be invaluable in investigating root causes. For example, if a worker was injured, you can review footage leading up to it: Was PPE worn? What was the chain of events? The key is that footage must be managed carefully – respecting privacy and not shared beyond necessary channels.

4.4 Privacy and Ethics Consideration

Collecting continuous video raises privacy considerations for workers. It is important to communicate to the workforce what is being monitored and why, and ensure it is done for legitimate safety and quality reasons. Some approaches include:

- Focus cameras on work zones, not break areas.
- Use AI to filter out irrelevant footage (so hours of people having lunch are not being stored).
- Possibly anonymize or not retain data longer than needed.
- Ensure compliance with any local laws on surveillance.

4.5 Example – Audit Trail for Concrete Pour

On a building project, the quality plan says no concrete pour without an inspection sign-off. With AI assistance: the site engineer runs the AI scan; the AI highlights any issues which are fixed; the system saves the annotated image of rebar and inserts it into the pour log; the engineer then signs off. If a later question arises about rebar adequacy, the saved evidence shows exactly what was in place before pouring. This audit trail (AI output + human approval) is stronger than either alone.

Knowledge Check

10. How does capturing AI-detected issues as images help in QA/QC audits?
11. Why is it important to store evidence of AI alerts even if they turned out to be false alarms?
12. What privacy considerations should be kept in mind when continuously recording a site for AI analysis?

Module 5: Risk Ranking and Prioritization with Vision Insights

Overview: With AI providing numerous observations, PEs and managers need to decide which issues to tackle first. This module integrates risk-based prioritization using a combination of AI data and traditional risk factors to focus efforts.

5.1 Risk = Likelihood × Consequence

Risk is considered as a combination of likelihood of an adverse event and the consequence if it occurs. AI vision systems contribute to assessing likelihood by identifying near misses or quality issues indicating something could go wrong. The engineer considers consequence: e.g., missing rebar in a critical column has high structural failure consequence, whereas a missing cosmetic feature is low consequence.

5.2 Using AI Data for Likelihood Estimation

AI systems provide objective indications of risk likelihood or trend:

- If the AI logs 10 near-misses of forklifts with pedestrians in a week in a certain zone, that indicates high likelihood of a serious incident if nothing changes.
- If the AI sees repeated QA issues (e.g., many instances of improper welding), it suggests higher likelihood of a section failing tests.
- If PPE compliance events are rare (one hard hat violation per month), the likelihood of an accident from that cause might be lower – but if an increasing trend is observed, risk is rising.

5.3 Consequence Factor (Human Judgment)

Consequence still needs human judgment. Engineers should classify each type of detected issue into a consequence category:

- Category I (High): Could cause loss of life or major structural failure (e.g., critical safety system compromised, severe collision).
- Category II (Medium): Could cause injury or significant rework/delay (e.g., moderate fall injury potential, structural deficiency requiring repair).
- Category III (Low): Minimal impact beyond small cost or no safety hazard (e.g., cosmetic defect, minor deviation from plan with negligible effect).

5.4 Risk Matrix Application & FMEA Example

Using the above, each detected issue can be plotted on a risk matrix (Likelihood vs Consequence). Those in the high-high area are top priority. A formalized FMEA approach is also helpful. The table below shows a sample FMEA for two site hazards, illustrating how AI affects the Detection (D) rating:

In this FMEA, Severity (S), Occurrence (O), and Detection (D) are each rated on a 1–10 scale; Detection uses an inverse scale where 1 indicates the failure is almost certain to be detected and 10 means it is nearly impossible to detect. The Risk Priority Number (RPN = S × O × D) ranges from 1 to 1,000, with higher values warranting priority action.

Hazard (Failure Mode)	Severity (S)	Occurrence (O)	Detection (D)	Risk Priority (S×O×D)	Controls & Actions (AI & Manual)
Forklift–pedestrian collision	10 (fatal)	4 (Occasional near-miss)	3 (AI monitors zone, but not everywhere)	120	Increase physical separation (barriers), assign spotters; continue AI monitoring; safety briefing for drivers.
Worker fall from height (no harness)	9 (severe injury)	2 (Rare)	2 (AI watches PPE continuously)	36	Ensure strict PPE enforcement; AI will catch if it happens, but maintain human oversight during critical activities.

This shows how AI affects the detection rating: in the forklift case, detection is moderate because AI covers some but not all site zones; in the harness case, detection is high (low D number) because AI is always watching for PPE compliance.

5.5 Using Priorities in Practice

- High risk issues (e.g., forklift near-miss cluster): Convene a safety meeting immediately, implement engineering controls (barriers or route changes) as top priority. Document in a safety bulletin.
- Medium risk issues (e.g., repeated moderate QA deviations): Schedule a meeting with the responsible crew to retrain and tighten procedures; add an extra QC check if needed.
- Low risk issues (e.g., sporadic minor non-compliances): Monitor via normal supervision; keep in a log for trend observation.

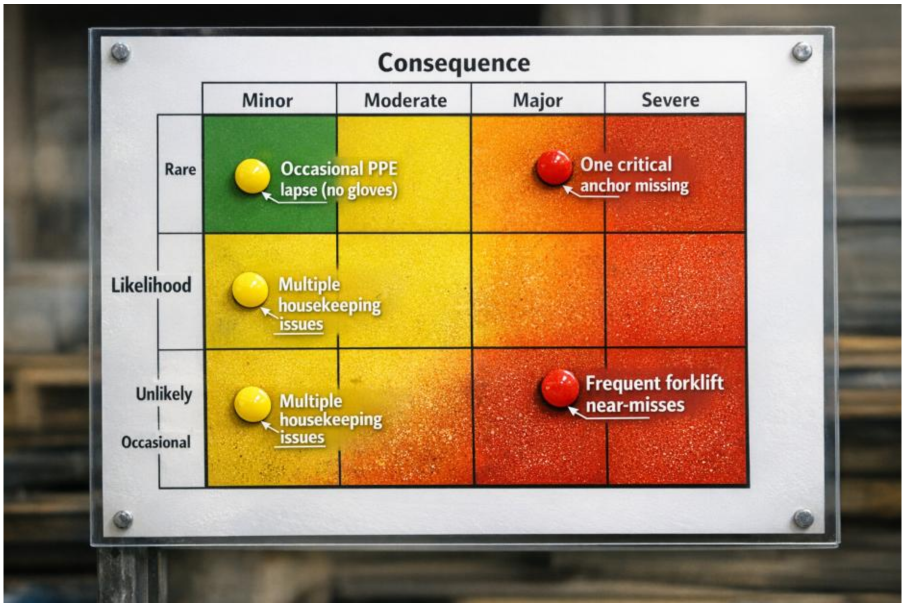

Knowledge Check

13. What two factors are considered when determining risk in a risk matrix?
14. If an AI system flags numerous low-severity issues and one high-severity issue, which should you address first and why?
15. In an FMEA, what does the Detection (D) rating represent in context of AI catching site issues?

Module 6: Human-in-the-Loop Workflows

Overview: AI outputs are only valuable if integrated into decision-making processes with humans in the loop at appropriate points. This module outlines how to incorporate AI vision into daily construction workflows, ensuring automated alerts are reviewed by people and then trigger real actions.

6.1 QA/QC Integration Example

On a building project, imagine a digital QA checklist enhanced by AI:

- Before concrete pour, the AI scans formwork and rebar. It produces a quick report: e.g., "All rebar present and within tolerance except possibly missing stirrup at grid C3." This appears on the site engineer's tablet.
- The site engineer goes to location C3, verifies a stirrup was indeed missing, and has the crew correct it immediately. The engineer then marks the checklist as OK and authorizes the pour.
- This process is formalized: company policy might state "AI-assisted pre-pour review required for all major pours, followed by engineer approval."
- The engineer signs off whether the AI finds issues or not – the AI speeds up finding issues, but final accountability rests with the human.

6.2 Safety Integration Example

For safety, consider a highway construction site with near-miss detection:

- Cameras monitor interactions between workers and moving equipment.
- The AI is configured to send an immediate alert (phone notification or site alarm) if a near-collision is detected. It logs the video snippet.
- The safety officer receives the alert and, if on site, intervenes right away. Otherwise, they investigate soon after by talking to the crew and reviewing the video.
- Each morning, the system compiles a safety report. The safety officer brings findings to the morning safety briefing for immediate corrective action.

6.3 Communication and Response Protocol

Integrating AI means establishing clear protocols:

- Who gets notified for each type of alert (QA manager for quality issues, safety manager for safety issues, site superintendent CC'd for anything affecting operations).
- How to respond: Standard operating procedures (SOPs) such as "If AI flags a structural QA issue, pause work in that area until a qualified engineer inspects and clears it."
- Escalation paths: If similar issues keep cropping up, escalate to higher management or specialists.

6.4 Integration with Existing Tools

Modern projects often use management platforms (Procore, Autodesk Construction Cloud (formerly BIM 360), etc.). The AI system's outputs can hook into these:

- If an AI finds a quality issue, it could automatically create a punch list item or Non-Conformance Report in the system, with the evidence attached and assigned to the responsible party.

For progress tracking, AI-measured completion percentages can feed into schedule tracking software or billing systems. One platform reported 3× faster draw processing using objective AI tracking of installed work *(per vendor data)*.[12]

6.5 Training and Acceptance

When introducing AI on site, involve the team:

- Explain to workers and supervisors what the AI is checking and why (it is to help ensure safety/quality, not to punish). This fosters acceptance and trust.
- Train personnel to handle AI alerts: the team must know how to interpret an alert and how to acknowledge it in the system once resolved.
- Maintain a feedback loop: if field personnel notice the AI often flags something incorrectly, they should inform the project engineers or AI vendors to improve the model.

6.6 Mechanical & Electrical Integration

Consider a mechanical/electrical scenario – an AI monitors an equipment yard for proper storage and lockout/tagout compliance. If it sees a generator running without a barricade or an electrical panel left open, it alerts the MEP supervisor. The site team then fixes it and logs that action. Similarly, AI could monitor if firestopping is installed around mechanical penetrations, and flag anything missed before walls are closed. The same human review principles apply across all disciplines.

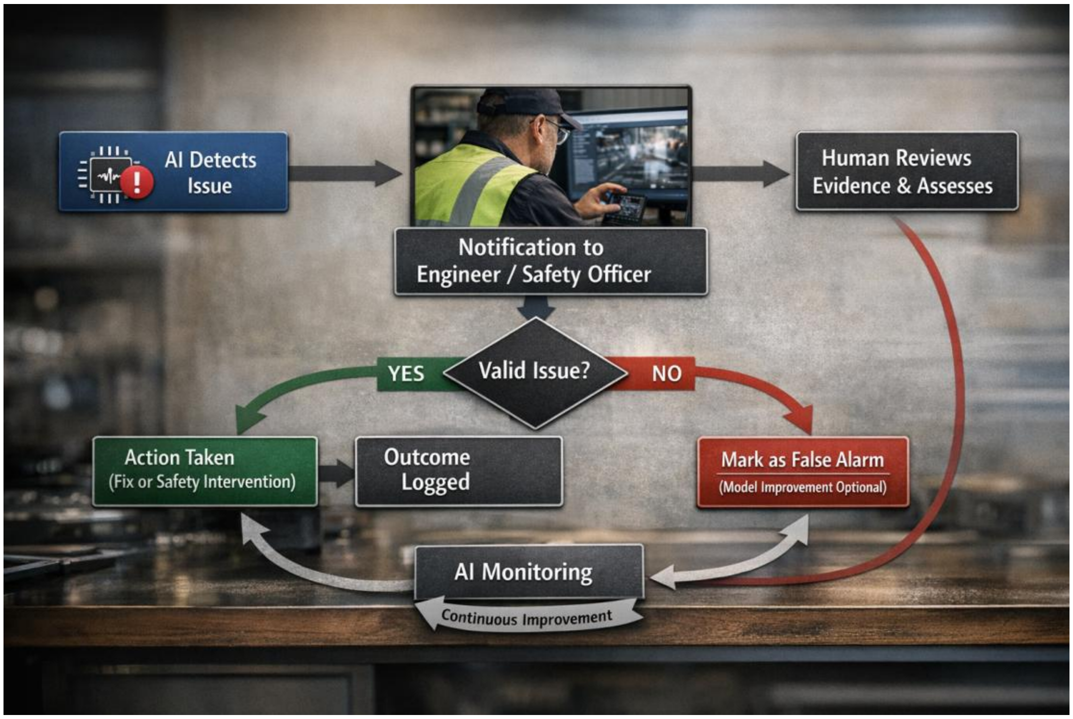

Knowledge Check

16. What is meant by a 'human-in-the-loop' in the context of AI for construction?
17. After an AI system flags an issue (say, a missing rebar), what should be the next step before taking further action?
18. How might AI-generated alerts be integrated into a project's existing quality or safety management tools?

Module 7: Documentation, Compliance, and Professional Responsibility

Overview: This final module addresses how to document and govern the use of AI in construction tasks to remain compliant with professional standards, quality management, and legal requirements.

7.1 Policy and Oversight

Firms implementing AI for QA/QC or safety should consider developing internal policies. Key elements include:

- Scope of AI Use: Document which activities AI assists and clarify that final determinations are by human inspectors. This could be in the project QA plan or safety plan.
- Responsibility Assignment: Define who oversees the AI outputs (e.g., QA/QC Manager reviews AI-generated inspection results daily; Site Safety Officer monitors and responds to AI safety alerts).
- Data Management: Outline how images/videos are stored, for how long, and how privacy is protected.
- Model Validation and Maintenance: Include a process to validate AI performance (perhaps by parallel testing on a pilot project first) and periodic re-evaluation (like monthly reviewing random samples of AI decisions).

7.2 ISO 9001 Alignment

Using AI can dovetail with a company's QMS:

- Risk-Based Thinking: Incorporate AI-derived insights into risk assessments – include near-miss analysis in the project's risk register to show proactive quality risk management.
- Continuous Improvement: Use metrics from the AI system – track "issues detected before failure" or reduction in rework thanks to AI, and report in management reviews.
- Internal Audits & Traceability: If ISO auditors ask, demonstrate that pre-pour inspections use an AI vision check per documented procedure, and that results are recorded as part of the inspection checklist.

7.3 Professional and Legal Defensibility

If something goes wrong, the presence of AI does not remove the engineer's accountability. Key principles:

- Always ensure a PE signs off on critical decisions. If using AI data in a final report or certification, treat it as you would any test result – verify and then sign with confidence.
- Document decisions: If AI flagged something and you decided it was not an issue, write down why, so if that choice is later questioned you can explain the rationale.
- Having the AI logs and your responses will help show that you were actively managing safety/quality, not negligent. Know the limits of your tools and plan around them.

7.4 Regulatory Compliance

Currently, no typical building code or OSHA rule specifically requires or regulates using AI on site – these are extra tools. But if you use them, ensure they complement rather than conflict with compliance obligations. OSHA construction standards (29 CFR 1926) require specific safety programs, hazard assessments, and documented controls depending on the work type — AI monitoring supplements but does not satisfy these obligations. Using AI does not replace that but can feed into it. Treat AI alerts with the seriousness they deserve and include them in your safety logs. Demonstrating an active hazard detection system can positively influence regulators and clients by showing you go beyond minimum compliance.

7.5 Ethical Use and Transparency

The ethical stance is clear: AI is a tool, and its benefits must be balanced with professional judgment. When writing proposals or communicating with clients, mention AI usage as an innovation but clarify that it operates under expert supervision. The NSPE encourages emerging tech like AI as long as PEs maintain responsible charge and ensure safety is paramount.

To maintain public trust and defensibility:

- Be transparent internally so everyone on the team knows how AI is being used and what it is not deciding by itself.
- Never ignore standard safety/quality practices thinking AI alone will catch something – use it to reinforce, not replace.
- Keep learning: as AI evolves, keep updated on best practices and any relevant guidelines that emerge.

Knowledge Check

19. What is one reason why thorough documentation of AI-assisted QA/QC activities is important for ISO 9001 compliance?
20. Who remains responsible if an AI fails to detect a construction flaw that later causes an issue?
21. How can engineers ensure that AI use on site is consistent with ethical obligations to protect public safety?

References

[1] U.S. Bureau of Labor Statistics Census of Fatal Occupational Injuries (CFOI).

[2] IJIRMPS (2024). Near-miss reporting statistics post AI safety system deployment. Ibid.

[3] IJIRMPS (2024). 34% reduction in total safety incidents over one year following CV-based PPE and hazard monitoring. Ibid.

[4] IJIRMPS (2024). YOLOv5 and deep-learning PPE detection model training. Ibid.

[5] IJIRMPS (2024). Object detection models (YOLO, Faster R-CNN) for construction site monitoring. Ibid.

[6] iFactory. "Construction Progress Monitoring Uses AI Vision." Available at: https://ifactoryapp.com/industries/infrastructure-management/construction-progress-monitoring-uses-ai-vision

[7] iFactory. Automatic counting of installed components for real-time quantity updates. Ibid.

[8] Intellisee. "Computer Vision: Occlusion, Low Light, and Adversarial Conditions." Available at: https://intellisee.com/intelligence/computer-vision-occlusion-low-light-adversarial-conditions/

[9] Intellisee. Detection error rates under heavy occlusion (5× or more accuracy degradation). Ibid.

[10] Intellisee. Overlapping camera placement recommended for critical zones. Ibid.

[12] iFactory. 3× faster draw processing using AI-tracked objective progress data. Op. cit.