Final Project Report

Industrial IoT System for Automated AI-Defect Detection of Concrete Sleepers

Administration	
Lead researcher	Dr. Felip Marti
Contact details	
University	Swinburne University of Technology
Date of submission	27 Feb 2026

Table of Contents

Acknowledgements and Project Team 4
Executive Summary / Abstract 6
Part A - Industry Problem, Needs and Relevance 7
Part B - Project Objectives, Hypotheses and Impact Pathways 10
Part C - Project Methodology and Key Activities 13
Part D - Analysis of Key Results, Deliverables, Description of Outputs and Project IP 17
Part E - Next Steps – Commercialisation & Impact Plan 41
Part F - Education and Training 46
Part G - Conclusions, Implications and Recommendations 47
Part H - Publication List 49
Part I - Appendices and Attachments 50

Acknowledgements and Project Team

Swinburne University of Technology gratefully acknowledges the financial support of **SmartCrete CRC**, which enabled the development and pilot deployment of this project.

We also sincerely thank **Sunset Concrete** for their collaboration, access to production facilities, and ongoing operational support throughout the installation, testing, and live validation of the system. Their partnership was essential to the project's success.

Project Team members and their roles in the project.

Title	Name	Position	Role in the Project	Organisation
Dr	Felip Marti	Senior Research Fellow	Project Lead	Swinburne University of Technology
Prof.	Prem Prakash Jayaraman	Professor	Provide consultation and lead the IoT-based components	Swinburne University of Technology
A/Prof.	Hailing Zhou	Associate Professor / AI Expert	Provide consultation and lead the AI-based algorithm development	Swinburne University of Technology
Prof.	Chris McCarthy	Associate Professor / Computer Vision Expert	Provide computer vision expertise	Swinburne University of Technology
Prof.	Pat Rajeev	Professor	Provide consultation and expertise on concrete mechanical properties and help with validation	Swinburne University of Technology
Prof.	Jay Sanjayan	Professor	Provide consultation and expertise on concrete mechanical properties and help with validation	Swinburne University of Technology
Mr	Vikram Sachdeva	Industry 4.0 Consultant	Providing manufacturing expertise and support to the Senior Research Fellow during project development	Swinburne University of Technology
Ms	Shihan Xu	Research Assistant	AI model development and defect detection algorithms	Swinburne University of Technology
Mr	Shuo Wang	Senior Technical Officer	Mechanical mounting, wiring, system commissioning, and compliance	Swinburne University of Technology
Mr	Josip Karabotic Milovac	Senior Research Engineer	Graphical user interface and database design	Swinburne University of Technology
Mr	Damian Blumenkranc	Project Steering Member	Providing strategic guidance, facilitating access to essential	Sunset Concrete

			facilities and data, and serving as the primary end-user	
Mr	Simon Thorne	COO / Technical Support	Providing strategic guidance, facilitating access to essential facilities and data	Sunset Concrete
Mr	Ruben Velasquez Ballenas	Technical Support	Providing metal fabrication capabilities and plant installations	Sunset Concrete

Executive Summary / Abstract

Precast concrete manufacturing relies heavily on manual quality control processes, which can be inconsistent, labour-intensive, and difficult to scale with increasing production speeds. Visual defect detection, dimensional checks, and reinforcement verification are typically performed under time pressure, limiting repeatability and traceability. These constraints can lead to rework, inefficiencies, and reduced data visibility for process improvement, affecting manufacturers and infrastructure reliability.

This project addressed these challenges through the development and live deployment of an AI- and IIoT-enabled automated quality control system at Sunset Concrete. The objectives were to determine whether machine vision and industrial sensors could operate at full production speed without disrupting workflow, achieve reliable defect and dimensional detection, and generate structured production data to enhance quality assurance.

The system was developed through staged milestones including KPI benchmarking, sensor selection, laboratory prototyping, AI model development, and industrial pilot testing. It has operated continuously in the plant since November 2025 without interrupting production. During validation, the vision system processed 377,164 frames in one month of live production. Pattern classification achieved 99.75% accuracy and colour classification 94.23%. Defect detection achieved high recall (≥0.93) for major defect classes and 100% recall for cracks. Height measurements across 12,898 sleepers demonstrated 98.8% compliance within ±2 mm tolerance under full production throughput. Rebar detection achieved a 0.09% error rate during stable calibration periods.

Width measurement variability highlighted the importance of mechanical stability in industrial AI deployments, but this was attributable to physical disturbances rather than computational limitations. Overall, the project successfully demonstrated that AI- and IIoT-based automated quality control can operate reliably in a live precast manufacturing environment, providing consistent inspection performance and structured digital traceability for the concrete sector.

Part A - Industry Problem, Needs and Relevance

1.1 Industry Context and Problem Definition

Sunset Concrete is undertaking a significant enhancement of its Quality Control (QC) processes to ensure that its concrete sleepers comply with applicable Australian Standards and relevant International Organization for Standardization (ISO) requirements.

At present, QC at the Officer manufacturing facility relies predominantly on manual visual inspection. While experienced operators are capable of identifying many surface and dimensional defects, this process is inherently time-consuming and susceptible to human error, fatigue and variability in judgement. These limitations introduce the risk of inconsistent defect detection, undetected non-conformances and variability in compliance outcomes. In a production environment where throughput, consistency and compliance are paramount, such variability represents a material operational risk.

The specific problems this research seeks to address are:

- Inconsistent and subjective defect detection due to manual inspection processes.
- Time-intensive QC procedures that constrain production efficiency.
- Limited capacity to systematically collect and analyse defect data for process improvement.
- Reactive, rather than preventative, quality management practices.

Collectively, these challenges restrict the organisation’s ability to optimise manufacturing performance, reduce waste and maintain continuous improvement aligned with national and international standards.

1.2 Barriers Preventing Current Resolution

Despite widespread recognition of the limitations of manual QC systems, several factors have prevented the industry from fully transitioning to automated, data-driven quality inspection.

First, the implementation of advanced automation—particularly Industrial Internet of Things (IIoT) and artificial intelligence (AI)-based inspection systems—has traditionally involved high upfront capital investment. Small-to-medium manufacturers may perceive such systems as financially prohibitive, particularly in industries with tight margins and fluctuating demand.

Secondly, technical complexity presents a significant barrier. Integrating AI-based visual inspection systems into existing production lines requires specialised expertise in machine vision, data analytics, systems integration and industrial networking. Many concrete manufacturers lack in-house digital engineering capability, and commercially available solutions may not be tailored to the specific geometries and defect types associated with precast concrete sleepers.

Thirdly, there is often limited empirical evidence within the Australian precast concrete sector demonstrating the operational performance, robustness and scalability of AI-driven QC systems under real manufacturing conditions. This uncertainty increases perceived risk and slows adoption.

Finally, traditional manufacturing cultures may favour established manual practices, particularly where these are perceived as reliable, even if inefficient. This can create organisational resistance to digital transformation initiatives.

1.3 Stakeholders Affected and Industry Costs

The challenges associated with manual QC processes are experienced directly by manufacturers such as Sunset Concrete, as well as indirectly by infrastructure clients and end users.

For manufacturers, undetected defects can result in costly rework, rejected batches, product returns and reputational risk. Rework not only increases direct labour and material costs, but also disrupts production schedules and reduces plant throughput. Where defective sleepers are identified post-delivery, transport and replacement costs further amplify financial impacts. While exact dollar values vary by defect rate and production scale, even marginal improvements in defect detection and prevention can translate into substantial cost savings over time in high-volume production environments.

Beyond direct financial costs, manual QC systems limit productivity by allocating skilled personnel to repetitive inspection tasks rather than higher-value activities such as process optimisation and preventive quality management. The absence of structured defect data also restricts evidence-based decision-making, constraining the organisation’s capacity to identify root causes and implement targeted improvements.

There are also environmental and sustainability implications. Defective products that require rework or disposal increase material waste, embodied energy consumption and associated carbon emissions. In an industry under growing pressure to demonstrate environmental responsibility, inefficient QC systems undermine broader sustainability objectives.

1.4 Research Solution and Anticipated Impact

This project proposes the development and pilot implementation of an Industrial Internet of Things (IIoT) system for automated, AI-based detection and reporting of defects in concrete sleepers. The system is designed to operate autonomically: automatically inspecting each sleeper, identifying defects and reporting results to operators in real time.

The research outcomes include:

a) Development of a proof-of-concept IIoT-enabled automated QC system for defect detection in concrete sleepers.
b) A four-week pilot trial at the manufacturing site to evaluate system performance, including detection accuracy, inspection speed, robustness and scalability.
c) A comprehensive final report detailing insights from system development and field testing, including analysis of collected defect data and recommendations for manufacturing process optimisation.

The anticipated impacts of successful adoption are multifaceted. Economically, automation of QC processes is expected to reduce error rates, lower rework and return costs, improve throughput and enhance labour productivity by reallocating human resources to higher-value tasks. Over time, the system’s data-driven insights may enable defect prevention at source, generating sustained cost efficiencies.

Socially, improved quality assurance enhances end-user confidence in the structural integrity, durability and safety of precast concrete products used in the built environment. Consistent compliance with Australian and ISO standards reduces the likelihood of premature failure, defects or performance variability in installed elements, thereby supporting public safety and long-term asset reliability. Enhanced and demonstrable quality control processes also strengthen industry credibility, fostering trust among clients, contractors and communities in the performance of precast concrete systems.

Environmentally, early and accurate defect detection reduces waste generation, minimises unnecessary material use and supports more sustainable manufacturing practices.

Importantly, while the immediate focus is on concrete sleepers, the insights derived from this project have broader relevance across the precast concrete sector. Many precast products face similar QC challenges, and a validated IIoT-based inspection framework could be adapted to other applications. As such, the research has the potential to catalyse wider industry transformation towards data-driven, automated quality management systems, supporting improved manufacturing standards and long-term competitiveness within the Australian construction materials industry.

Part B - Project Objectives, Hypotheses and Impact Pathways

2.1 Project Objectives and Targeted Industry Challenges

This project directly addresses the limitations of manual quality control (QC) processes in precast concrete manufacturing, particularly those relating to inspection speed, consistency, data utilisation and integration with modern production systems. Specifically, the project responds to the following industry challenges:

- The inability of manual inspection processes to consistently match the pace of automated production lines.
- Variability in defect detection outcomes due to human subjectivity and fatigue.
- Limited structured data collection to support predictive or preventative quality management.
- Low levels of digital integration between inspection processes and broader manufacturing systems.

The primary objective of the project is to design, develop and pilot a proof-of-concept Industrial Internet of Things (IIoT)-enabled automated QC system capable of detecting and reporting defects in concrete sleepers in real time. A further objective is to evaluate the operational feasibility, performance reliability and scalability of the system within a live production environment.

Through a four-week pilot deployment at Sunset Concrete, the project seeks to assess inspection accuracy, processing speed, system robustness and integration capability. An additional objective is to generate and analyse structured defect data to inform future manufacturing process optimisation.

2.2 Key Hypotheses and Technical Assumptions

The project investigates several interrelated hypotheses concerning the technical and operational feasibility of automated QC in precast concrete production.

The first key hypothesis is that the automated IIoT system will be capable of identifying defects at a speed compatible with the existing production line, without creating bottlenecks or disrupting workflow. Success in this respect is defined by the system’s ability to inspect each sleeper within the available cycle time while maintaining reliable performance.

The second hypothesis is that the data collected from labelled sleeper samples will be sufficient in quality and quantity to train an artificial intelligence (AI) model capable of achieving acceptable detection accuracy. This includes the assumption that defect types can be consistently labelled and categorised, and that the dataset will support model generalisation to new production batches. Success will be measured through quantitative performance indicators such as detection accuracy, false positive and false negative rates, and overall model robustness under operational conditions.

The third hypothesis is that integration of the IIoT system into the QC workflow will increase overall process efficiency by reducing human error and reallocating human operators to higher-value tasks. Here, success is defined not only in technical terms, but also in operational and organisational terms—specifically, measurable reductions in inspection time variability, improved defect traceability, and positive operator engagement with the system.

Several technical assumptions underpin these hypotheses. It is assumed that the selected combination of sensors, cameras and lighting configurations will produce imaging quality sufficient for accurate defect detection and measurement. The reliability of image capture under varying environmental conditions (e.g. dust, vibration, lighting variability) is critical to system success.

It is also assumed that human operators will engage constructively with the digital transformation introduced by the system. Operator acceptance and effective use of system outputs are essential to realising productivity and decision-making benefits.

2.3 Intended Impacts and Value Proposition

The project aims to unlock measurable economic, environmental and operational value within the precast concrete industry.

From an economic perspective, successful implementation of an automated QC system has the potential to:

- Reduce defect-related rework and associated labour and material costs.
- Lower the incidence of product returns and non-conformance events.
- Increase throughput by aligning inspection speed with production rates.
- Improve labour productivity through redeployment of skilled personnel to process improvement and supervision roles.

Even modest reductions in defect rates across high-volume production can generate substantial cumulative savings. Additionally, structured defect data enables root cause analysis and preventative action, supporting sustained cost reductions over time.

From an environmental standpoint, improved early-stage defect detection reduces material waste, embodied energy loss and unnecessary disposal of non-conforming products. By shifting from reactive correction to preventative quality management, the system supports more resource-efficient manufacturing practices aligned with broader sustainability objectives in the construction materials sector.

For end users, the value proposition lies in improved product consistency, enhanced reliability and demonstrable compliance with relevant standards. For the industry partner, the business case includes improved operational efficiency, enhanced competitiveness and the potential to position itself as a leader in digitally enabled manufacturing.

The project aligns with broader governmental and institutional priorities relating to Industry 4.0 adoption, advanced manufacturing transformation, digital innovation and sustainable production practices within Australia. By supporting the integration of IIoT and AI technologies into traditional manufacturing contexts, the project contributes to national objectives aimed at strengthening productivity, technological capability and industrial resilience.

2.4 Impact Pathways and Adoption Strategy

The immediate pathway to impact involves pilot implementation and performance validation within the manufacturing operations of Sunset Concrete. The four-week trial will generate empirical evidence regarding technical feasibility, operational integration and measurable benefits. This evidence base will inform refinement of the system design and development of an internal business case for continued deployment.

In the short term, adoption will focus on incremental scaling within the existing facility, including system optimisation, operator training and integration with existing production management systems. Project participants will play a central role in evaluating performance data, refining AI models, and documenting operational improvements to support broader uptake.

In the medium term, successful validation may support expansion across additional production lines or product categories within the organisation. The structured data generated during ongoing operation can be leveraged to continuously improve both the AI model and manufacturing processes.

In the longer term, the proof-of-concept and pilot findings may inform commercialisation pathways, including adaptation of the IIoT-enabled QC framework for other precast concrete applications facing similar inspection challenges. Industry-wide dissemination of lessons learned (through technical reporting, case studies and collaborative engagement) can support broader adoption of automated, data-driven QC systems across the concrete manufacturing sector.

Through this staged implementation pathway, the project seeks to move from technical validation to operational integration and ultimately to scalable industry transformation.

Part C- Project Methodology and Key Activities

3.1 Methodological Framework and Alignment with Hypotheses

The project methodology was structured to systematically test the technical and operational hypotheses outlined in Section 2, namely: (i) that an automated IIoT system can operate at production-line speed; (ii) that sufficient, well-labelled data can be collected to train an accurate AI-based defect detection model; and (iii) that integration of the system can improve QC efficiency and reliability.

The methodology comprised six sequential and interrelated phases: (1) KPI establishment and defect definition; (2) hardware and algorithm selection; (3) laboratory-based prototype development; (4) on-site data collection; (5) AI model development and integration; and (6) pilot testing in a live production environment. At each phase, measurable performance criteria were defined to determine whether assumptions were validated or required refinement.

3.2 Establishment of KPIs and Development of a Defect Dictionary

The first phase focused on defining performance benchmarks and defect scope in collaboration with Sunset Concrete.

KPI Establishment. Existing QC performance indicators were reviewed, including inspection time per unit, rejection rates, rework frequency and operator workload. In parallel, AI-specific metrics were introduced to stakeholders, including precision, recall, F1-score, false positive rate and false negative rate.

Defect Dictionary Development. A structured defect dictionary was developed in consultation with operational and quality personnel. This documented and defined key defect categories, including dimensional variances, surface cracks, bubbles (size and density), chips (size and density), discolouration, stains, and missing reinforcement (including rebar depth considerations).

This structured taxonomy provided the foundation for consistent data labelling and AI model training, directly addressing the hypothesis concerning the sufficiency and quality of labelled data. A subset of priority defects was selected for inclusion within the initial IIoT system scope to ensure feasibility and focus during prototyping.

3.3 Sensor Selection, Desktop Research and System Architecture Design

The second phase involved hardware and software selection supported by targeted desktop research.

Hardware and Sensor Selection. The project team identified and specified:

- High-resolution industrial cameras for visual defect detection.
- Laser-based measurement sensors for dimensional verification.
- Rebar detection and depth measurement sensors for concrete cover assessment.
- Lighting systems to ensure consistent imaging conditions.
- Compute and storage hardware to support real-time data processing.

The assumption that appropriate sensor combinations could deliver reliable imaging and measurement performance was tested through specification review, controlled laboratory trials and vendor benchmarking.

Desktop Research on ML/DL Approaches. A structured review of contemporary machine learning (ML) and deep learning (DL) techniques for concrete defect detection was undertaken. Convolutional neural networks (CNNs) and related computer vision architectures were evaluated for suitability in identifying surface anomalies and dimensional inconsistencies. The selection of candidate models was informed by reported performance in peer-reviewed literature and industrial case studies involving precast or concrete-based materials.

3.4 Laboratory-Based Prototype Development and Testing

A laboratory-scale IIoT prototype was designed and developed at the Swinburne University of Technology Factory of the Future facility.

The prototype incorporated selected sensors, mounting systems and a tunnel-like inspection enclosure designed to simulate production-line conditions. Concrete sleeper samples supplied by Sunset Concrete were used for controlled experimentation.

Testing activities included:

- Image quality validation under controlled lighting variations.
- Sensor calibration and dimensional measurement accuracy testing.
- Evaluation of rebar detection sensors for depth and cover performance.
- System throughput testing to simulate production cycle times.

Performance was measured against predefined thresholds for imaging clarity, measurement tolerance and processing speed. Any deviations informed iterative refinement of hardware configuration and enclosure design.

3.5 On-Site Data Collection and AI Model Development

Following laboratory validation, the prototype was deployed on-site for structured data collection within the manufacturing environment.

Data Capture and Preparation. Real-time images and sensor readings were collected as sleepers passed through the inspection zone. Data preparation included cleaning, formatting and structured labelling based on the previously defined defect dictionary. Labelling was conducted in collaboration with QC personnel to ensure alignment with operational standards.

The sufficiency of the dataset was evaluated in terms of:

- Number of samples per defect category.
- Class balance across defect and non-defect instances.
- Inter-rater consistency in labelling.

Algorithm Development and Testing. AI-based detection algorithms were developed and iteratively refined. Model performance was evaluated using quantitative metrics including:

- Precision and recall for each defect category.
- Overall classification accuracy.
- False positive and false negative rates.
- Processing latency per unit inspected.

Integration testing assessed the combined performance of hardware and software, verifying that detection outputs were delivered within the allowable production cycle time.

3.6 Pilot Deployment and Operational Validation

A four-week pilot trial was conducted at Sunset Concrete under full production conditions.

The objectives of the pilot were to validate:

- Real-time inspection capability without production disruption.
- System robustness under environmental variability (e.g. dust, vibration).
- Detection performance relative to human inspectors.

Continuous performance monitoring was undertaken, and operator feedback was collected to evaluate system desirability and integration feasibility.

3.7 Prior Work and Knowledge Foundations

The project built upon established research in machine vision, AI-based defect detection and Industry 4.0-enabled manufacturing systems. Prior academic and industry literature in computer vision for construction materials and automated industrial inspection informed the selection of sensing configurations, data processing pipelines and candidate machine learning architectures.

No background intellectual property (IP) was transferred or licensed as part of this project. Rather, the work drew upon publicly available research, open-source machine learning frameworks and established engineering design principles. The technical contribution lay in the contextual adaptation, integration and validation of these approaches within a precast concrete manufacturing environment.

Industry knowledge was contributed through operational expertise provided by Sunset Concrete, including detailed understanding of production workflows, defect definitions, quality thresholds and inspection practices. This domain-specific knowledge was essential in shaping the defect dictionary, defining performance criteria and ensuring alignment between algorithm outputs and real-world quality requirements.

The integration of contemporary AI methodologies with site-specific manufacturing knowledge enabled the development of a tailored, application-specific solution. The novelty of the project therefore resides not in proprietary IP, but in the practical translation of established digital technologies into a live precast concrete production setting, supported by empirical validation under operational conditions.

3.8 Novelty and Competitive Positioning

While commercial machine vision inspection systems exist globally, many are designed for high-uniformity manufacturing sectors (e.g. automotive or electronics) and are not tailored to the material variability and surface complexity inherent in precast concrete products.

The novelty of this research lies in:

- The integration of multimodal sensing (visual, dimensional and rebar detection) within a unified IIoT framework.
- The development of a context-specific defect dictionary aligned with Australian manufacturing practices.
- The coupling of AI-driven inspection with structured data analytics for process optimisation, rather than solely pass/fail classification.
- Empirical validation within an operational precast manufacturing environment.

This project advances beyond generic off-the-shelf solutions by embedding domain-specific knowledge, enabling adaptability to product variability and supporting long-term manufacturing optimisation.

3.9 Industry Engagement and Feedback Mechanisms

Throughout project delivery, engagement with industry stakeholders was continuous and structured. This included:

- Joint KPI definition workshops.
- Collaborative defect dictionary development.
- On-site data collection and validation sessions.
- Operator training and AI literacy discussions.
- Feedback sessions during pilot deployment.

This participatory approach ensured that system development remained aligned with operational realities, regulatory requirements and user expectations. It also supported organisational readiness for digital transformation and strengthened pathways to adoption beyond the pilot phase.

Through this rigorous, staged methodology, the project systematically tested its hypotheses, validated technical feasibility and demonstrated operational viability within a real-world manufacturing context.

Part D- Analysis of Key Results, Deliverables, Description of Outputs and Project IP

This section synthesises the technical results obtained across the major system components, aligns them with the project hypotheses and milestones, and evaluates how the outcomes address the industry challenges identified. Detailed technical results are presented in Section 4.1, 4.2 and 4.3. The discussion below provides an integrated narrative of findings and implications.

4.1 Machine Vision Outcomes

The vision module was designed to detect concrete sleepers, identify sleeper attributes (e.g. pattern type and colour), and classify selected surface defects in accordance with the agreed defect dictionary. To understand the performance of the vision model, this section will illustrate its performance in terms of each function: defect detection, sleeper pattern identification, and sleeper colour identification.

To evaluate the vision model's performance, we utilized images and logs saved during December 2025. This evaluation dataset was collected over a 31-day period with continuous monitoring of the concrete sleeper inspection process, representing real-world data from an industrial production environment.

During December 2025, the vision system processed 377,164 total frames, of which 26,454 frames (approximately 7%) were saved for further analysis due to the detection of defects or anomalies. The following evaluation is based on these saved images and the system-provided logs.

Tabel 1. Vision model and evaluation dataset

Model version	**V1.0**
Duration of Collection	December 2025 (December 1-31, 2025)
Total frames processed	377,164
Frames saved with defect detected	26,454

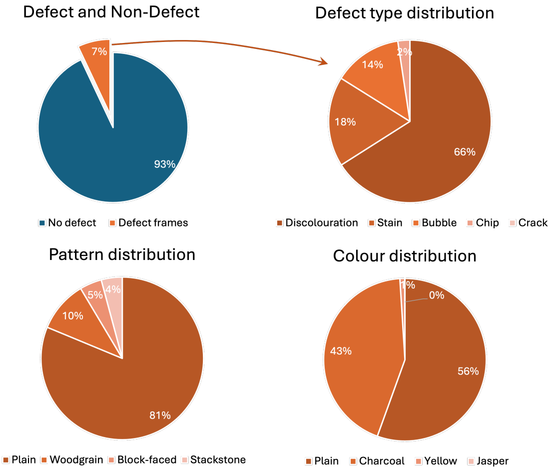

Figure 1. Distribution analysis regarding to defect types, pattern types and sleeper colour

The Plain pattern dominates the dataset with over 80% of occurrences, which is consistent with standard concrete sleeper production. Woodgrain represents the second most common pattern at 10.19%, while Block-faced and Stackstone patterns occur less frequently, together comprising less than 9% of the dataset.

The dataset is primarily composed of Plain (55.52%) and Charcoal (43.48%) coloured sleepers, which together account for nearly 99% of all colour classifications. Yellow sleepers are present but rare at less than 1%, while Jasper-coloured sleepers are only 0.04% of the dataset. Discolorations are the most frequently detected defect type, accounting for 65.91% of all defect detections (17,437 frames). This is followed by Stains (17.88%, 4,730 frames) and Bubble defects (13.68%, 3,620 frames). Chip defects are less common, representing 2.40% of defects, while Crack defects are extremely rare at only 0.12% (31 occurrences).

4.1.1 Introduction of Evaluation Metrics[1]

Precision is the proportion of all the model's positive classifications that are positive. It is mathematically defined as:

$$Precision = \frac{correctly\ classified\ actual\ positives}{everything\ classified\ as\ possitive} = \frac{TP}{TP + FP}$$

- Precision tells us, out of all sleepers the system labelled as a given class, how many really were that class?
- Example: if the system says 100 sleepers have Stains and 80 of them truly have Stains, then precision for Stains is 80%.
- High precision means few false alarms for that class.

Recall is also known as the true positive rate (TPR)**,** or the proportion of all actual positives that were classified correctly as positives.

$$Recall(or\ TPR) = \frac{correctly\ classified\ actual\ positives}{all\ actual\ positives} = \frac{TP}{TP + FN}$$

- Recall tells us, out of all sleepers that really have a given class, how many did the system correctly flag?
- Example: if there are 100 sleepers with Cracks in the dataset and the system detects 95 of them, recall for Cracks is 95%.
- High recall means few missed defects.

Accuracy is the proportion of all classifications that were correct, whether positive or negative. It is mathematically defined as

$$Accuracy = \frac{correctly\ classifications}{total\ classification} = \frac{TP + TN}{TP + TN + FP + FN}$$

- Accuracy tells us, out of all sleepers in the evaluation set, how many did the system classify into the correct colour / pattern / defect class?
- Example: if there are 400 sleepers and the system gets 310 completely right (colour or defect type), overall accuracy is 310 / 400 = 77.5%.

F1-score is the harmonic mean (a kind of average) of precision and recall.

$$F1 = 2 * \frac{precision * recall}{precision + recall} = \frac{2TP}{2TP + FP + FN}$$

A high F1-score for a defect class (for example, bubble defect) means that most true bubble defects are detected (good recall) and most alarms are correct (good precision).

In the following sections, we will use these metrics to evaluate the model performance on colour classification, pattern classification and defect type detection.

---

[1] https://developers.google.com/machine-learning/crash-course/classification/accuracy-precision-recall#:~:text=Exercise%3A%20Check%20your%20understanding&text=Exercise%3A%20Check%20your%20understanding

4.1.2 Evaluation of Concrete Sleeper Defect Type Classification

For defect classification, the system achieves an overall accuracy of 77.26% across six classes (Bubble, Chip, Crack, Discoloration, Stain, NoDefect). As shown in the confusion matrix (Figure 2) and metric table (Table 2), Bubble, Chip, Discoloration and Stain all have high recall (≥ 0.93) and good precision (around 0.76–0.86). This means that when a sleeper truly contains one of these defect types, the system usually detects it, and most of the corresponding alarms are correct. Examples of correct and incorrect defect classification results are shown in Figure 3.

Crack defects maintain 100% recall but have low precision (0.16), meaning the system tends to over-report crack defects rather than miss any. This lower precision is primarily due to extremely limited crack training data, which affects the model's ability to distinguish crack defects from similar visual patterns (e.g., scratches and line-shaped stains). Despite this limitation, the high recall ensures that all actual crack defects are detected, which is important for safety given that crack defects represent serious structural damage. In practice, this trade-off has limited operational impact: crack detections are extremely rare, accounting for only 0.12% of all defect detections (31 frames out of 26,454 defect-detected frames in the December sample), so the low precision does not generate a significant number of false alarms for manual inspection.

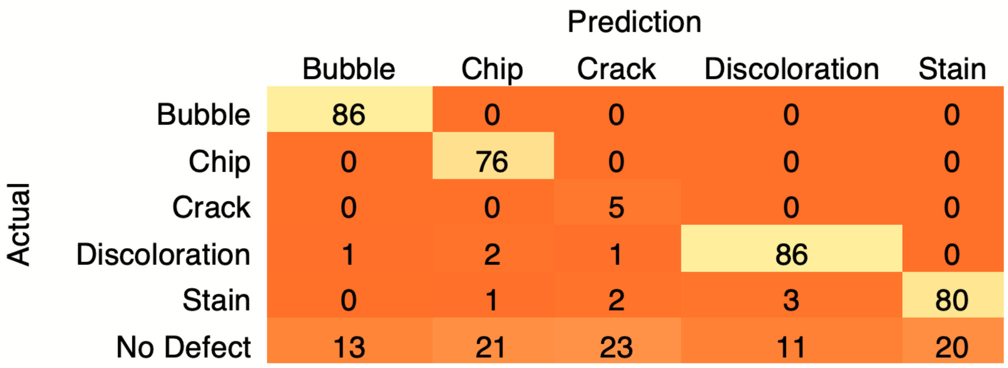

		Prediction				
		Bubble	Chip	Crack	Discoloration	Stain
Actual	Bubble	86	0	0	0	0
	Chip	0	76	0	0	0
	Crack	0	0	5	0	0
	Discoloration	1	2	1	86	0
	Stain	0	1	2	3	80
	No Defect	13	21	23	11	20

Figure 2. Confusion matrix for defect type classification

Table 2. Performance metrics for defect type classification

	Precision	Recall	F1-score
Bubble	0.86	1	0.92
Chip	0.76	1	0.86
Crack	0.16	1	0.28
Discoloration	0.86	0.96	0.91
Stain	0.8	0.93	0.86
Overall accuracy			77.26%

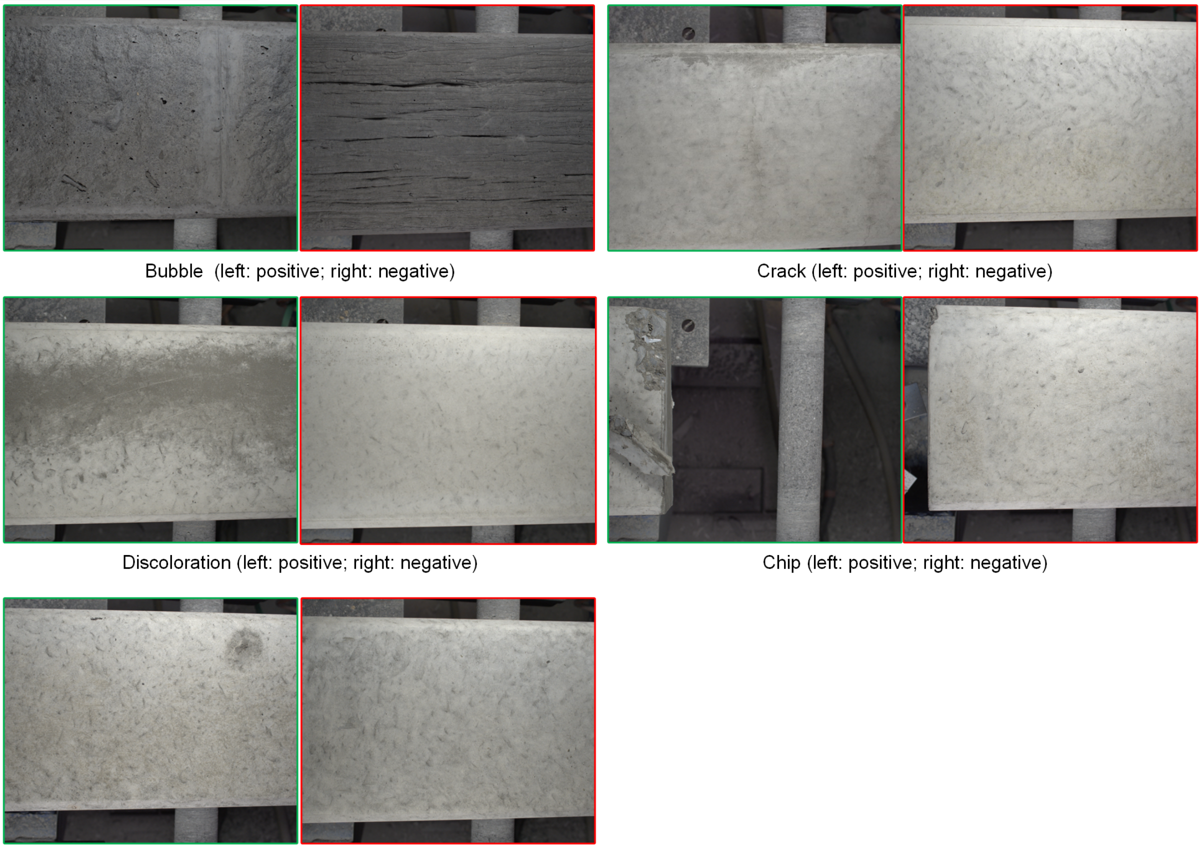

Bubble (left: positive; right: negative)

Crack (left: positive; right: negative)

Discoloration (left: positive; right: negative)

Chip (left: positive; right: negative)

Stain (left: positive; right: negative)

Figure 3. Example defect-classification results. Green-bordered images show correct colour predictions. Red-bordered images show false predictions.

4.1.3 Evaluation of Concrete Sleeper Pattern Type Classification

The model delivers excellent performance on pattern classification, achieving 99.75% accuracy across the four texture classes (Block-faced, Smooth, Stackstone, Woodgrain). Block-faced and Woodgrain sleepers are classified perfectly, with precision and recall both equal to 1.00. There is a single misclassification where a frame of a Smooth sleeper edge is predicted as Stackstone (top-right image in Figure 5), which slightly reduces the precision and recall for Smooth and Stackstone respectively; in this case the edge region provides insufficient visual cues for the model. Figure 4 and Table 3 show the per-class performance of the pattern classifier, summarising precision, recall, and F1-score for each pattern type, and overall demonstrate that the system performs consistently and accurately on pattern classification.

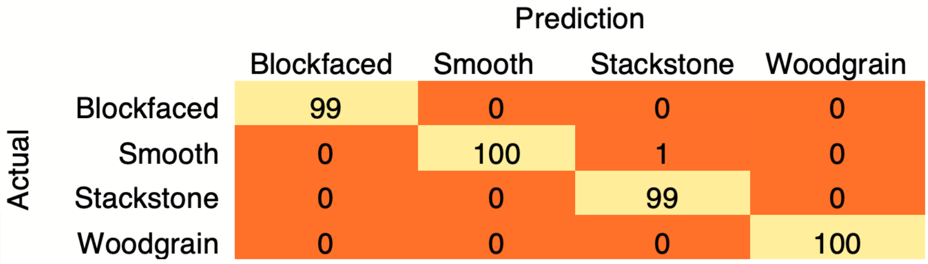

Actual \ Prediction	Blockfaced	Smooth	Stackstone	Woodgrain
Blockfaced	99	0	0	0
Smooth	0	100	1	0
Stackstone	0	0	99	0
Woodgrain	0	0	0	100

Figure 4. Confusion matrix for pattern type classification

Table 3. Performance metrics for pattern type classification

	Precision	Recall	F1-score
Block-faced	1	1	1
Smooth	1	0.99	1
Stackstone	0.99	1	0.99
Woodgrain	1	1	1
Overall accuracy			99.75%

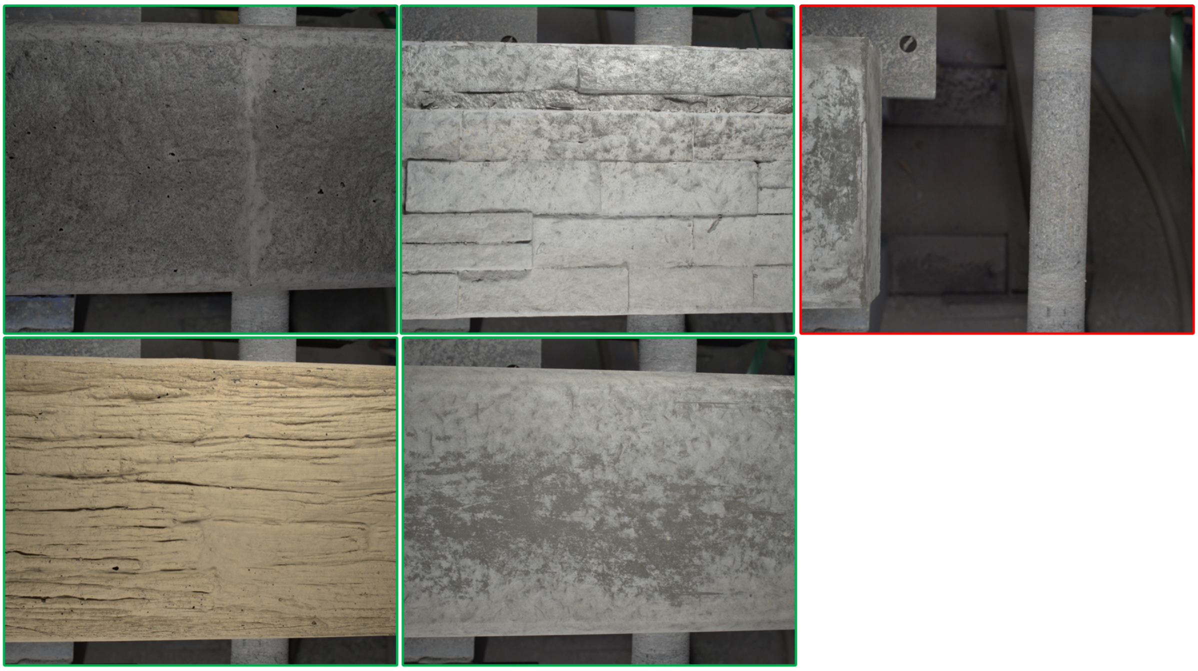

Figure 5. Example pattern-classification outputs. Green-bordered images show correct colour predictions. Red-bordered images show false predictions. (top-right) a frame of Smooth sleeper edge falsely predicted as Stack-stone.

4.1.4 Evaluation of Concrete Sleeper Colour Type Classification

The colour classification algorithm demonstrates strong performance across the four sleeper colours (Charcoal, Jasper, Plain, Yellow), achieving 94.23% accuracy on the matched evaluation set. Figure 6 and Table 4 together illustrate the per-class performance of the colour classification model, summarising the precision, recall, and F1-score for each colour type. Yellow and Jasper are detected with perfect recall, indicating the model reliably recognises these colour categories when present. The major failure cases occur when Plain sleepers are falsely predicted as Charcoal (21 Plain sleepers misclassified as Charcoal). This error is primarily driven by surface discolorations on otherwise Plain sleepers, where darker defect regions visually resemble the darker tone of Charcoal and mislead the model's colour decision. In addition, Background frames (no sleeper surface present) reveal a separate error type: false positive sleeper detection, where the model assigns a sleeper colour to a non-sleeper scene. Figure 7 summarises these behaviours, showing representative correct predictions as well as typical false-positive and colour-confusion cases.

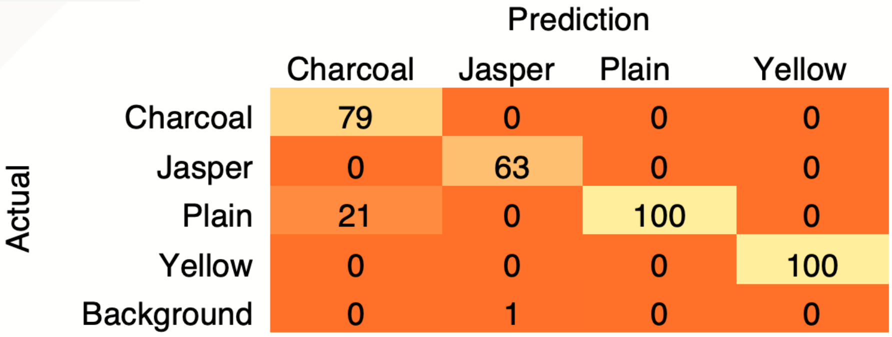

Actual \ Prediction	Charcoal	Jasper	Plain	Yellow
Charcoal	79	0	0	0
Jasper	0	63	0	0
Plain	21	0	100	0
Yellow	0	0	0	100
Background	0	1	0	0

Figure 6. Confusion matrix for colour type classification

Table 4. Performance metrics for colour type classification

	Precision	**Recall**	**F1-score**
Charcoal	0.79	1	0.88
Jasper	0.98	1	0.99
Plain	1	0.83	0.9
Yellow	1	1	1
Overall accuracy			94.23%

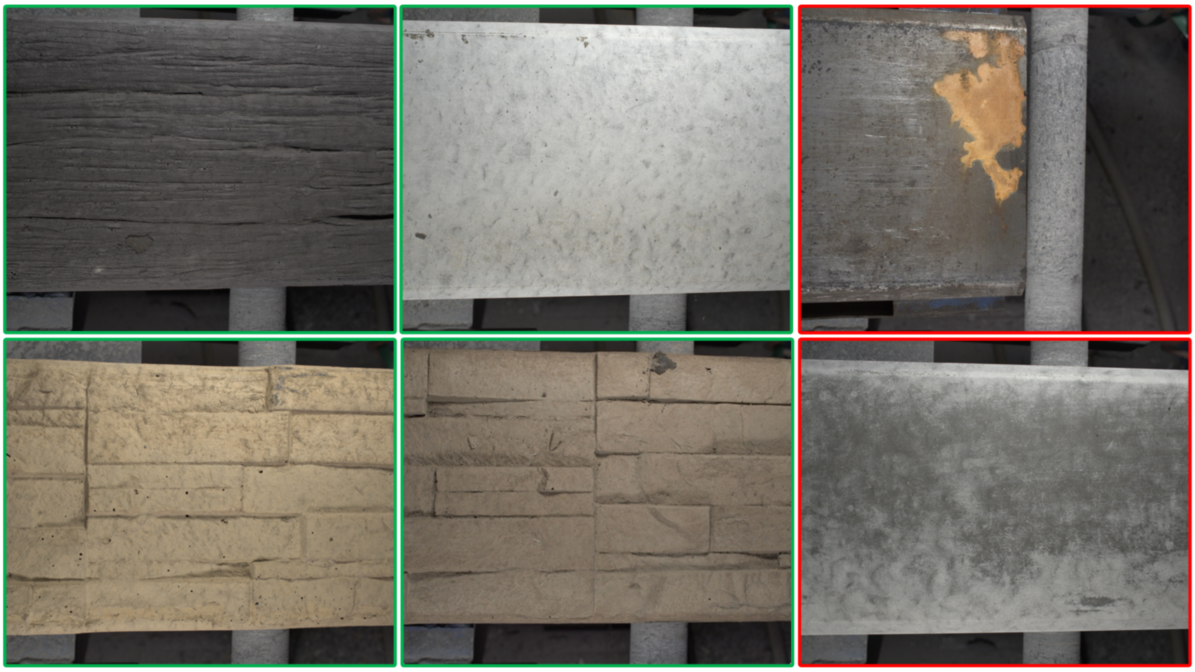

Figure 7. Example colour-classification outputs. Green-bordered images show correct colour predictions. Red-bordered images show false predictions. (top-right) a Background frame with no sleeper surface present that is falsely detected as a Jasper sleeper; (bottom-right) a Plain sleeper with Discolorations that is falsely detected as Charcoal.

Overall, the evaluation shows that the vision module can reliably support automated inspection of concrete sleepers in a real factory environment. Pattern classification is highly accurate (99.75%), meaning the system can consistently distinguish among different sleeper patterns. Colour classification also performs strongly (94.23% accuracy), with errors largely limited to visually ambiguous cases where discolouration defects darken Plain sleepers. Defect-type classification achieves 77.26% accuracy with high recall for the main defect categories (Bubble, Chip, Discolouration, Stain) and deliberately conservative crack detection, so serious damage is unlikely to be missed. These results support the project’s key assumption that a deep-learning–based vision system can provide robust and accurate inspection performance on real production data, reducing reliance on manual visual checks while still allowing operators to review and refine edge cases.

4.1.5 Key Results and Implications

The evaluation of the vision module during the production period of December 2025, comprises 377,164 processed frames and 26,454 defect-flagged frames, demonstrates that AI-based inspection is technically viable in a live precast manufacturing environment.

The results provide strong empirical support for the project’s core hypotheses.

Alignment with Production-Line Requirements
The system operated continuously across the evaluation period without interrupting the production workflow, confirming that inference speed and system latency are compatible with production-line cycle times. This directly validates the hypothesis that an automated IIoT vision system can operate at required industrial throughput without introducing bottlenecks.

Data Sufficiency and Model Robustness
The scale of the evaluation dataset (collected under real operating conditions) demonstrates that sufficient labelled data can be captured within a production environment to train and validate deep learning models.

Pattern classification achieved 99.75% accuracy across four texture classes, indicating near-perfect reliability in distinguishing sleeper surface types. Colour classification achieved 94.23% overall accuracy, with misclassifications largely attributable to visually ambiguous cases (e.g., discolouration on Plain sleepers).

For defect detection, the system achieved 77.26% overall accuracy across six classes. Critically, recall for the primary defect categories (Bubble, Chip, Discolouration, Stain) was ≥0.93, indicating that the system successfully identifies the vast majority of true defects in these categories. This is particularly important from a quality assurance perspective, where missed defects (false negatives) pose greater operational risk than false positives.

Conservative Crack Detection Strategy
Crack detection achieved 100% recall, ensuring that no true crack defects were missed during evaluation. Although precision for cracks was low (0.16), this is primarily attributable to the extremely small number of crack instances (31 frames, representing 0.12% of defect detections). From an operational risk perspective, this represents a conservative detection strategy (prioritising defect capture over minimising false alarms. Given the low base rate of crack occurrence, the impact of lower precision on workflow remains limited.

Operational Implications
The high recall rates for major defect categories and near-perfect pattern classification indicate that the system can reliably support automated or semi-automated inspection workflows. The structured logging of defect type, pattern and colour provides traceable digital records, directly addressing the industry challenge of limited QC data capture identified in Section 1.

Importantly, performance characteristics indicate that the system can:

- Reduce reliance on purely manual visual inspection;

- Improve repeatability and consistency of defect identification;
- Support operator review of edge cases rather than full manual inspection of all units;
- Enable longitudinal defect trend analysis for process optimisation.

Overall, the results confirm the project's key assumption that a deep learning–based machine vision system can deliver robust inspection performance using real production data. While further dataset expansion (particularly for rare defect types) would likely improve class balance and precision metrics, the current performance demonstrates operational feasibility and provides a strong foundation for progressive automation of quality control processes.

4.2 Dimensional and Sensor-Based Measurement Results

In addition to surface defect detection, the system incorporates laser-based sensors for precise dimensional measurements, including both width and height of the concrete sleepers. These measurements, captured in real-time as the sleepers pass through the sensor zone, are critical for ensuring compliance with industry standards and quality controls. This section outlines the challenges faced in conducting dimensional measurements, particularly focusing on the width and height sensors, and presents the results of these measurements under operational conditions.

The width and height sensors work together to detect the presence of sleepers and accurately measure their dimensions. As the sleeper passes through the inspection frame, the sensors record data that is subsequently analysed to determine the width, height, and presence of any internal rebar. Given the longstanding use of photoelectric sensors for such tasks, the dimensional measurement system was integrated to meet the specific requirements of the industry partner. However, as this technology is not novel, the focus of this report is on evaluating the operational performance of the system within the context of the trial period and its alignment with existing quality control thresholds.

4.2.1 Challenges and Limitations

A significant challenge encountered during the trial period was the operational disturbances caused by the interaction between the sleepers and the measurement sensors. Specifically, sleepers were frequently lifted by operators—likely unintentionally—causing them to come into direct contact with the width sensors. These interactions resulted in the physical displacement or complete detachment of the sensor mounts. Unfortunately, these incidents were not always reported to the research team, which hindered the proper evaluation of the system's performance.

The displacement or removal of the sensors without recalibration severely impacted the accuracy of the measurements. Without proper recalibration or repositioning, the sensors could not accurately detect the dimensional data, resulting in unreliable width measurements. In addition, a small vertical step between the demoulding area and the sensor zone caused further challenges. This step led to slight misalignments as the sleeper transitioned from one zone to another, occasionally affecting the width readings at the leading and trailing edges of the sleeper.

Despite these challenges, the height measurements remained stable and reliable. The height sensors, positioned in a more controlled and stable part of the inspection line, were less affected by the sensor misalignment issues. The height measurements exhibited only minor fluctuations, with values consistently aligning closely with the nominal dimensions of the sleepers, indicating that the system performed well under the conditions of controlled sensor alignment.

4.2.2 Key Results and Observations

Width Measurements:
The width measurements, though captured accurately at the system's core, displayed significant fluctuations. This is exemplified by the data shown in the graph below, which presents the width measurements of the first 10 sleepers processed on 24 February 2026, following the system recalibration performed the previous day.

Although the first 10 sleepers measured between 07:02 AM and 07:08 AM are known to be smooth and nominally 100 mm wide, the recorded data from both sensors shows significant short-term fluctuations within each measurement window. Instead of a tight distribution around 100 mm, the readings frequently vary by ±3–5 mm, with occasional extremes approaching ±10 mm. Such variation is not physically realistic for smooth sleepers and therefore does not represent true dimensional changes. The similar scatter patterns observed in both Sensor 1 and Sensor 2 further suggest that the fluctuations are not due to isolated sensor calibration issues.

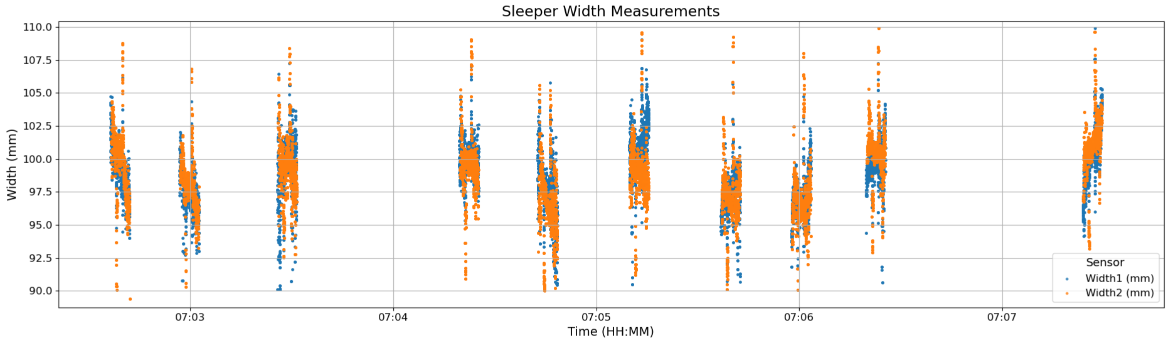

In the graph below, height measurements for the same sleepers were presented, showing much less scatter, with measurements remaining within a ±1–2 mm range of the nominal value. The improved stability of the height measurements suggests that the observed instability in the width measurements is primarily due to dynamic movement and sensor alignment issues rather than flaws in the measurement technology itself.

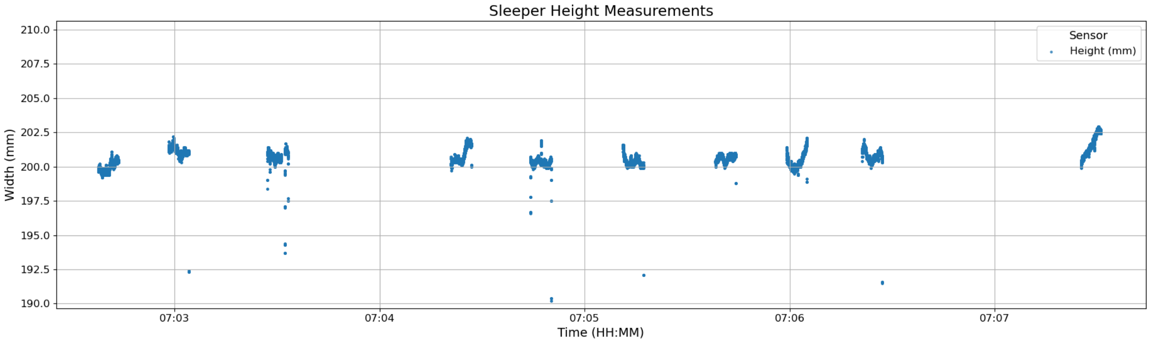

A key reason for this improved stability is the location of the height sensor. The height measurement system is positioned in a more stable area of the line, where the sleeper movement is better controlled and there are fewer disturbances in the measurement axis. Unlike the width station, it is not directly affected by the small step between the demoulding area and the sensor zone. As a result, there is significantly less bouncing or dynamic movement in the vertical measurement direction, leading to more consistent and reliable readings.

In the graph below presents the width measurements for the same 10 sleepers, but limited to only the first 10 data points (leading edge) and the last 10 data points (trailing edge) recorded for each sleeper. These are the specific values used by the system to estimate the final sleeper width. As in Figure 1, both Sensor 1 (blue) and Sensor 2 (orange) are shown, and each cluster corresponds to a single sleeper between 07:02 AM and 07:08 AM. Compared to the full raw dataset, the number of points per sleeper is significantly reduced, focusing only on the edge measurements that define the calculated width.

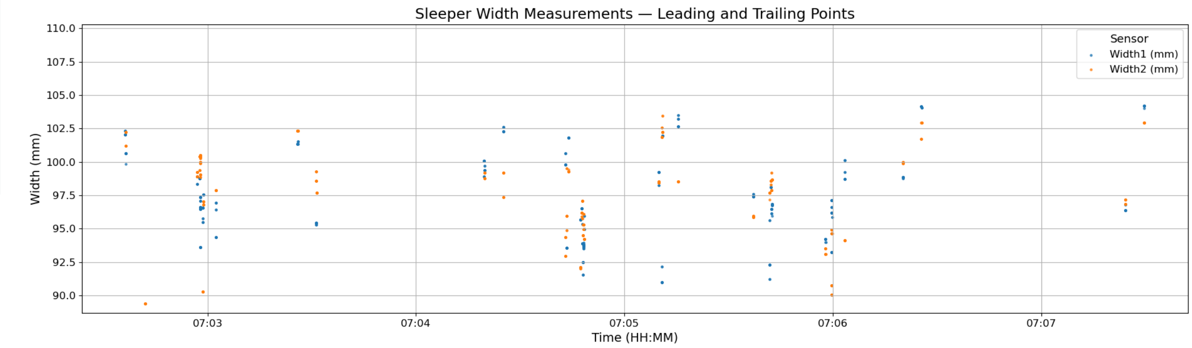

Although the fluctuations in this figure are not always as extreme as those observed in the complete raw signal (first graph in Section 4.2.2), the variability in the first and last 10 data points remains noticeable. Since these edge points are directly used for width estimation, any instability at the moment the sleeper enters or exits the sensing zone has a disproportionate impact on the final result. The scatter visible in several sleepers—particularly differences between leading and trailing values and inconsistencies between sensors—indicates that the measurement is still influenced by movement or bouncing. Therefore, even if the variation is not always extreme, it remains sufficiently inconsistent to question the reliability of the width estimation when based solely on these limited edge data points.

Based on the observed behaviour of the width measurements, the primary recommendation is mechanical stabilisation of the sleepers before and during the measurement process. The data strongly suggests that the variability is driven by dynamic movement (bouncing and vertical disturbance), particularly at the transition between the demoulding area and the sensor zone. Since the fluctuation is physical in nature, any attempt to correct it purely through software filtering or signal smoothing will only mask the symptom rather than eliminate the root cause. Software compensation may reduce noise visually, but it cannot reliably reconstruct the true geometric width if the sleeper position relative to the sensors is unstable at the time of measurement.

Height Measurements:

The height measurements performed during the trial demonstrated a high level of accuracy under controlled conditions. On 19 November 2025, a dedicated experimental setup was conducted in which 41 sleepers were passed through the system specifically for dimensional validation. These sleepers were secondments that had already been demoulded and were intentionally selected to avoid disrupting standard factory operations during working hours. This controlled session allowed the research team to perform comprehensive dimensional verification without interfering with production output.

The trial conditions differed from routine production flow. Because the sleepers had already been demoulded, it was operationally more challenging for plant personnel to manually feed them through the inspection frame. As a result, sleepers were passed individually, with sufficient spacing between each unit to ensure clean detection events and stable measurement acquisition. This spacing eliminated the risk of overlapping triggers or rapid consecutive passes, thereby creating an optimal measurement environment.

All 41 sleepers were measured for both width and height. The sample set included sleepers from different mould patterns and varying nominal widths, ensuring that the dimensional validation was not restricted to a single geometry. For height specifically, every automated measurement recorded by the system matched the corresponding manual measurement obtained by the research team, resulting in a 100% agreement rate for this dataset.

While this outcome demonstrates the technical capability and intrinsic accuracy of the height sensing system, it is important to contextualise the result. The exceptional accuracy achieved during this session is likely influenced by the highly controlled experimental conditions, including individual sleeper handling, reduced production pressure, and the absence of mechanical disturbances. Therefore, although the result confirms that the sensor technology is capable of precise height measurement, its performance under continuous high-volume production conditions must be evaluated separately, as discussed in the subsequent sections.

For the period between 2 February 2026 and 20 February 2026, a comprehensive dataset was collected following the mechanical reinforcement and physical protection of the height sensor assembly. These dates were deliberately selected for analysis, as they correspond to the interval during which the sensor mounts had been strengthened and shielded against accidental impacts from sleepers and plant handling equipment. This ensured that the data reflected performance under mechanically stable and properly aligned conditions.

During this period, a total of 12,898 sleepers were measured under live production conditions. Daily production volumes ranged from 528 to 1,091 sleepers, demonstrating that the sensor system was operating under realistic, high-throughput industrial demand rather than controlled laboratory conditions. The accompanying height distribution table presents the number of sleepers recorded at each discrete millimetre increment between 196 mm and 204 mm for each production day.

Day	Sleeper Height Size (mm)									Total per day
	196	197	198	199	200	201	202	203	204	
2026-02-02	0	4	13	101	382	171	25	0	0	696
2026-02-03	0	4	27	149	395	230	32	1	0	838
2026-02-04	0	2	10	51	334	401	118	7	0	923
2026-02-05	0	1	7	55	363	305	75	2	0	808
2026-02-06	1	4	8	62	248	178	24	3	0	528
2026-02-09	0	5	12	67	383	347	71	4	0	889
2026-02-10	1	6	23	107	472	373	85	3	0	1070
2026-02-11	2	3	15	56	387	499	120	9	0	1091
2026-02-12	1	6	32	130	442	240	53	5	0	909
2026-02-13	2	15	39	121	500	235	47	2	0	961
2026-02-16	2	3	7	67	358	377	109	9	0	932
2026-02-17	0	4	13	65	280	337	115	16	2	832
2026-02-18	2	6	46	170	425	164	32	5	0	850
2026-02-19	0	9	31	164	381	142	29	2	0	758
2026-02-20	0	6	28	117	420	196	44	2	0	813
Total per size	11	78	311	1482	5770	4195	979	70	2	12898

The aggregated results show a tightly centred distribution around the nominal design height of 200 mm. Of the 12,898 sleepers measured:

- 5,770 sleepers (44.7%) were recorded at exactly 200 mm;

- 4,195 sleepers (32.5%) were recorded at 201 mm;
- 1,482 sleepers (11.5%) were recorded at 199 mm.

Collectively, 11,447 sleepers—representing approximately 88.7% of total production during this interval—fell within the narrow band of 199–201 mm. This range lies entirely within the standard tolerance of ±1 mm from nominal and comfortably within the broader acceptable tolerance band of ±1–2 mm, which is configurable within the system parameters.

When the acceptable interval is extended to 198–202 mm, the compliance rate increases further to 12,737 sleepers, representing 98.8% of total production. Only a very small fraction of measurements fell outside this extended tolerance band. Specifically, 11 sleepers were recorded at 196 mm and 2 sleepers at 204 mm, together accounting for less than 0.1% of the total dataset. These extreme values are statistically insignificant in the context of nearly 13,000 measurements and may reflect isolated production variability or rare measurement edge cases rather than systematic sensor error.

Importantly, the consistency of the distribution across multiple consecutive production days further supports the conclusion that the height sensor maintained calibration stability over time. The modal value remained consistently at 200 mm or 201 mm across all reporting dates, and no progressive drift or widening of the distribution was observed. This temporal stability indicates that, once mechanically reinforced and protected from physical disturbance, the height measurement subsystem operated reliably under sustained industrial conditions.

Overall, the February dataset provides strong empirical evidence of measurement repeatability, accuracy, and robustness at production scale. Unlike the controlled November validation trial, this dataset reflects continuous plant operation, variable production speeds, and normal handling conditions. The sustained concentration of values within the nominal tolerance band demonstrates that the height sensing system is suitable for long-term deployment in an industrial environment, provided that mechanical protection and alignment integrity are maintained.

4.2.3 Implications for System Optimisation

The discrepancy observed in the width measurements highlights the need for mechanical improvements in the measurement zone. These improvements include:

- **Eliminating the Step**: The small vertical step between the demoulding area and sensor zone introduces unnecessary disturbances that affect measurement consistency. A ramped or levelled transition would help mitigate this issue.
- **Stabilising Sleeper Movement**: To reduce vibration and movement, lateral constraint rollers or shock absorbers could be installed to dampen dynamic motion before the sleeper enters the sensor zone.
- **Sensor Relocation**: Where possible, repositioning the width sensors to a more stable location with reduced vibration would also contribute to improved measurement accuracy.

While software-based signal smoothing may reduce the visual noise in the data, it is not sufficient to address the root causes of the measurement instability. As such, mechanical stabilisation and sensor realignment are essential for achieving more reliable and consistent width measurements.

4.2.4 Conclusion

The dimensional measurement system, once appropriately calibrated and mechanically stabilised, provides highly accurate height measurements, meeting the industry's stringent quality standards. The width measurement system, however, continues to exhibit significant variability due to dynamic movement, vibrations, and sensor misalignment. Although the data shows the system's potential, especially in controlled conditions, further optimisation of the measurement zone is required to improve the consistency and reliability of width measurements. Addressing these issues will enhance the overall performance of the system and facilitate more reliable automated quality control for sleeper production.

4.3 Rebar Detection Assessment

4.3.1 Trial Period and Operational Constraints

The rebar detection subsystem was designed to operate in conjunction with the width sensors, which trigger the detection event when a sleeper passes through the inspection frame. Accurate alignment and calibration of both the width sensors and rebar sensors are therefore critical to reliable operation.

For the purpose of this evaluation, the data analysed comes from log files generated by the system during the trial period. These logs contain detailed records, including timestamped events, sensor readings, and detection results. The data includes pass/fail results for both Rebar 1 (R1) and Rebar 2 (R2), as well as the raw sensor values that inform these decisions. This data is crucial for assessing system accuracy and troubleshooting operational anomalies.

Although the broader machine vision system operated throughout December 2025, the valid evaluation window for rebar detection was restricted to the period between **9 December and 18 December 2025**. The sensors were recalibrated and confirmed to be correctly aligned on 9 December. On 18 December at approximately 7:49 am, abnormal detection behaviour was observed, indicating probable physical misalignment of one or more sensors. No formal notification of mechanical disturbance was received from site personnel.

Subsequent investigation strongly suggests that physical interference with the inspection frame occurred during plant operations. Operators lifting sleepers contacted structural components of the inspection rig, including the frame, lighting assemblies and sensor mounts. In at least one documented instance, a width sensor was found physically displaced and resting on top of the frame rather than aligned toward the sleeper path. As the rebar detection algorithm relies on accurate sleeper presence triggering from the width sensors, this misalignment rendered subsequent rebar data unreliable.

These operational disturbances significantly constrained the duration of valid rebar evaluation and highlight an important implementation consideration: industrial inspection hardware must be robustly protected against unintended plant interaction, and operational communication protocols must be established to report sensor displacement events immediately.

4.3.2 False Triggers and Error Clarification

Some false triggers were recorded during the evaluation period. These events occurred when the width sensor laser beams were interrupted by external objects (such as operator activity), causing the system to falsely trigger a rebar evaluation. Importantly, these false triggers are not counted in the final error analysis since they do not reflect actual detection failures, but rather environmental disruptions (e.g., operator hands during cleaning). The focus remains on genuine errors where rebar detection was either missed or inconsistent.

4.3.3 Performance During Stable Calibration Period

During the stable and calibrated period (9–17 December), the rebar detection system performed reliably relative to production volume. Here's a summary of the performance on monitored days:

Date	Sleepers Demoulded	Errors (Valid Detections)
10 December	964	0 errors (8 false triggers)
11 December	39	0 errors
12 December	518	0 errors
15 December	869	0 errors (6 false triggers)
16 December	978	0 errors
17 December	994	4 errors (3 false triggers)

Performance Analysis:

- **Total Sleepers Processed** (9–17 December): **4362 sleepers**
- **Total Valid Errors**: **4 errors**
- **False Triggers**: **9 false triggers** (these are excluded from final error calculations)
- **Error Rate**: **0.09%** (4 valid errors / 4362 sleepers processed)

During the stable alignment period, the system successfully processed thousands of sleepers with **minimal confirmed detection errors**. Excluding false triggers, the error rate remained exceptionally low, demonstrating strong performance.

On **17 December**, **4 potential detection issues** were logged where **Rebar 1** detected rebar, but **Rebar 2** did not. These cases are likely due to:

- Marginal threshold sensitivity
- Minor sensor displacement

Importantly, these cases were isolated and represent a very small proportion of total inspected sleepers. Threshold sensitivity can be adjusted in software to increase detection responsiveness, suggesting these events are tunable system parameters rather than structural performance limitations.

4.3.4 Sensor Misalignment Event on 18 December

On 18 December, at approximately 7:49 am, 60 fail events were recorded within a two-minute window. This production rate is physically impossible and indicates a significant disturbance to the width trigger sensor.

Following this event, rebar detection behaviour became unstable, with increased false triggers and inconsistent readings. Given the physical dependency between width detection and rebar evaluation, misalignment of the trigger sensor invalidates subsequent rebar data.

Because no formal incident report was provided at the time of impact, diagnosis relied solely on log interpretation and anomaly detection. This underscores a key implementation lesson: industrial AI inspection systems require not only technical robustness but also procedural alignment with plant operations.

Data collected after this event were therefore excluded from formal performance evaluation.

4.3.5 Implications for Hypothesis Validation

Despite operational disruptions, the results during the calibrated window support the project's technical assumptions:

1. **Sensor Feasibility:** The selected rebar sensing technology is capable of detecting rebar presence and depth under live production conditions.
2. **Integration Capability:** The rebar system successfully integrated with the IIoT architecture and logging framework.
3. **Threshold Adjustability:** Detection sensitivity can be tuned through software parameters, enabling mitigation of marginal detection cases.

The principal limitations observed were not algorithmic or sensing failures, but mechanical robustness and operational integration challenges. These are engineering and procedural issues rather than fundamental technical barriers.

4.3.6 Key Lessons and Recommendations

The trial highlighted several critical requirements for future deployment:

- Reinforced and protected sensor mounting assemblies to prevent displacement.
- Physical shielding of laser trigger systems from incidental contact.
- Clear site communication protocols for reporting hardware disturbance.

Addressing these factors will significantly improve long-term reliability.

4.3.7 Overall Assessment

When evaluated during periods of confirmed calibration and stable hardware positioning, the rebar detection subsystem processed thousands of sleepers with minimal confirmed detection anomalies.

The majority of logged “failures” were attributable to:

- Laser beam interruption without sleeper presence, or
- Mechanical disturbance of sensor alignment.

The underlying sensing technology demonstrated strong potential for reliable automated concrete cover assessment.

Accordingly, the rebar module is considered technically viable, with future performance improvements dependent primarily on mechanical hardening and operational integration rather than algorithmic redevelopment.

4.3.8 Mechanical Robustness and Production-Line Integration Lessons

The trial period provided several important practical insights regarding the mechanical robustness of inspection hardware and the operational integration of automated quality control systems within a live manufacturing environment.

While the sensing technologies themselves demonstrated strong technical capability, the deployment highlighted that mechanical protection and plant interaction are critical determinants of system reliability. During routine operations, sleepers are frequently handled manually by plant personnel immediately after demoulding. In several instances during the trial, sleepers were unintentionally lifted into contact with the inspection frame or sensor assemblies. These interactions caused physical displacement of sensor mounts, particularly for the width sensors and, indirectly, the rebar detection trigger system that relies on the width sensors for event initiation.

Because these disturbances were not always reported to the research team immediately, the sensors occasionally remained misaligned for extended periods before recalibration occurred. As demonstrated in Sections 4.2 and 4.3, even small positional shifts can significantly affect measurement accuracy or trigger behaviour. The experience therefore highlights an important deployment requirement: inspection systems operating in high-throughput industrial environments must be designed with robust mechanical protection and tolerance to incidental operator contact.

Future installations should therefore prioritise:

- Reinforced structural mounting of sensors and frames
- Physical shielding of sensitive laser measurement components
- Mechanical design that prevents direct contact between sleepers and sensing hardware
- Clear operational protocols for immediately reporting and recalibrating any sensor displacement

Another practical consideration identified during the trial relates to the handling of non-compliant products detected by the automated system. The prototype installation focused primarily on detection and data collection rather than automated product diversion. Consequently, when defects or dimensional anomalies were identified, the system generated alerts and logged the events, but the physical removal of the corresponding sleeper from the production flow remained a manual process performed by plant operators.

While this approach was appropriate for the pilot phase, a fully integrated production deployment would ideally include a formal rejection or isolation mechanism, such as:

- a flagged inspection output communicated to operators via the HMI,
- a marking or tagging mechanism applied to non-compliant sleepers, or
- integration with downstream handling equipment to divert flagged products to a quarantine area.

Implementing such a mechanism would allow the inspection system to move beyond monitoring and become a fully closed-loop quality control tool within the production process.

Overall, the trial demonstrated that the primary challenges encountered were not related to sensing technology or AI performance, but rather to the mechanical and operational realities of installing precision inspection equipment within an active industrial environment. These insights represent a valuable outcome of the pilot deployment and provide clear guidance for improving robustness, maintainability, and workflow integration in future implementations.

4.4 Validation Against Key Hypotheses

This section evaluates the extent to which the experimental and operational results presented in Sections 4.1 - 4.3 validate the key hypotheses and technical assumptions defined in Section 2.2. Overall, the findings provide strong support for the technical feasibility of an IIoT-enabled automated QC system in a live precast concrete manufacturing environment, while also identifying clear areas for mechanical and procedural optimisation.

Hypothesis 1: Compatibility with Production-Line Speed

The first hypothesis set out to validate how the automated IIoT system would be capable of operating at a speed compatible with the existing production line without introducing bottlenecks.

This hypothesis is **validated with strong empirical support**.

Since November 2025, the IIoT inspection system has operated continuously in the production environment without interrupting plant operations. The December 2025 dataset (377,164 processed frames over 31 days) represents the formal validation and analysis window rather than the beginning of system deployment. Throughout this period and beyond, inference speed, data processing, and logging architecture remained fully compatible with industrial cycle times, confirming that the system can run in parallel with live manufacturing without introducing bottlenecks. Similarly, during the formally analysed stable calibration window for rebar detection (9–17 December), the system processed 4,362 sleepers with an error rate of only 0.09%, again without disrupting production flow.

Dimensional measurements were likewise captured in real time under full production throughput. The February 2026 dataset—comprising 12,898 height measurements across consecutive high-volume production days—represents the structured analysis period following mechanical reinforcement of the sensor assembly. However, the dimensional subsystem had already been operating prior to this interval as part of the live installation. These results demonstrate that the sensing, processing, and IIoT data infrastructure can scale to sustained industrial demand.

Where performance limitations were observed (e.g., width instability or the rebar anomaly event on 18 December), these were attributable to mechanical disturbance and sensor alignment issues rather than computational latency, inference limitations, or system speed. From a systems engineering perspective, the hypothesis regarding compatibility with continuous industrial production is therefore confirmed.

Hypothesis 2: Data Sufficiency and AI Detection Performance

The second hypothesis set out to validate how labelled production data would be sufficient to train AI models capable of achieving acceptable detection accuracy and generalisation.

This hypothesis is **validated with strong empirical support**.

The December dataset provided a substantial real-world evaluation base, including 26,454 defect-flagged frames. Model performance demonstrated:

- **Pattern classification accuracy of 99.75%**, effectively near-perfect.
- **Colour classification accuracy of 94.23%**, with errors primarily in visually ambiguous cases.
- **Defect classification accuracy of 77.26%**, with high recall (≥0.93) for major defect types (Bubble, Chip, Discoloration, Stain).
- **Crack detection recall of 100%**, ensuring no critical structural defects were missed, albeit with conservative precision due to class imbalance.

These results confirm that sufficient real production data can be captured, labelled, and used to train robust models. Importantly, the system prioritised high recall for safety-critical defects, which aligns with industrial QC risk priorities.

A refinement to the original hypothesis may be considered: while data quantity proved sufficient overall, rare defect categories (e.g., cracks) would benefit from further dataset expansion to improve precision without compromising recall. This does not invalidate the hypothesis but identifies an optimisation pathway.

Hypothesis 3: Workflow Efficiency and Reduction of Human Variability

The third hypothesis set out to validate how integration of the IIoT system would improve QC efficiency by reducing human subjectivity and enabling higher-value operator tasks.

Below we provide a list of outcomes post deployment and pilot of the IIoT system that **validates this hypothesis**.

The system generates structured, timestamped digital records of defect type, pattern, colour, dimensional measurements, and rebar status. This directly addresses the previously identified industry challenges of limited data capture and poor traceability. The high recall rates for primary defects demonstrate improved repeatability compared to purely manual inspection, which is inherently subject to fatigue and variability.

In the case of height measurements, the February dataset shows exceptional consistency under production conditions (98.8% within ±2 mm tolerance), confirming that automated dimensional control can reliably replace manual gauge-based checks when mechanical stability is ensured.

However, operational disturbances—such as unreported sensor displacement—highlight that organisational alignment and operator engagement are critical to realising full efficiency gains. The technology itself proved capable; the remaining barriers are procedural and mechanical integration challenges rather than conceptual or algorithmic limitations.

Baseline Comparison and Measured Impact

A direct quantitative comparison against historical manual inspection performance was not fully available, as structured defect logging, inspection time tracking, and operator-level variability metrics were not formally recorded prior to system deployment. This reflects a broader industry challenge, where quality control processes are often manual, experience-based, and inconsistently documented.

However, the absence of a formal baseline itself highlights a key outcome of the project: the establishment of a measurable and repeatable quality framework.

Prior to system deployment:

- Inspection coverage was **non-uniform and operator-dependent**, with defects often identified either downstream in the process or at the customer end
- No complete dataset of defect occurrence existed
- Defect classification lacked standardisation and relied heavily on individual judgement
- Real-time feedback to production was limited, increasing the risk of defect propagation

Following deployment of the IIoT system, the QC process demonstrates measurable improvements:

- **100% inspection coverage** of all sleepers passing through the conveyor, compared to previously arbitrary or sample-based checks
- **Real-time defect detection**, reducing decision latency and enabling immediate corrective action
- **Standardised defect classification**, supported by a formalised defect dictionary developed during the project
- **Complete digital traceability**, with all defects timestamped, categorised, and stored for analysis
- **Reduced operator cognitive load**, with the system assisting across multiple defect types simultaneously

A representative one-week production dataset (9–13 February 2026) illustrates the scale and consistency of detection now achieved:

- Total sleepers inspected: **4,911**
- Defects detected:
    - Bubbles: 190
    - Discoloration: 253
    - Stains: 80
    - Chips: 60
    - Cracks: 3
    - Height deviations: 64

This level of defect visibility was not previously achievable under manual inspection conditions. In particular, the consistent detection of high-frequency, low-severity defects (e.g., bubbles, discoloration) demonstrates the system's ability to capture quality trends that would typically go unrecorded.

In addition, the system enables:

- Trend analysis across production days, shifts, and product types

- Early identification of process drift
- Data-driven decision-making for process improvement

While precise numerical comparisons against historical human error rates and inspection times are not available, the transition from **non-instrumented, subjective inspection** to a **fully digitised, real-time, and comprehensive QC system** represents a substantial step-change in capability.

4.5 Validation of Technical Assumptions

Several technical assumptions were also tested:

Sensor and Imaging Suitability
The vision system operated effectively under real factory conditions (dust, vibration, lighting variability), supporting the assumption that the selected hardware configuration can produce sufficient image quality for AI-based inspection.

Mechanical Stability Requirement
Results from the width sensor demonstrate that while the sensing technology itself is capable, physical instability (e.g., step transition, operator contact) directly affects measurement reliability. The contrast between unstable width data and highly stable height data confirms that mechanical robustness is a prerequisite for dimensional reliability. This insight strengthens, rather than weakens, the original assumptions by clearly identifying implementation conditions.

System Integration Feasibility
All modules (vision, dimensional sensing, rebar detection) successfully integrated within a unified IIoT logging framework. Data were timestamped, stored, and retrievable for analysis, confirming digital integration feasibility.

Operator Engagement Assumption
While no formal resistance to the system was observed, the impact of undocumented sensor disturbances suggests that clearer communication protocols and physical protection mechanisms are required. Future deployments should incorporate formal reporting procedures and hardware safeguarding to ensure sustained performance.

Overall Summary of Testing and Validation

The collected evidence across vision inspection, dimensional measurement, and rebar detection demonstrates that the proposed IIoT-enabled automated QC system is technically feasible, production-compatible, and capable of delivering high inspection accuracy under real industrial conditions.

The most significant limitations identified during the pilot were operational challenges that impacted mechanical robustness (such as sleeper hitting the frame) —not algorithmic inadequacy, data insufficiency, or computational performance. We included several mechanisms to address this issue such as reinforcing sensor and protecting them (as seen in the February height dataset), having protection for the camera etc. These improved overall mechanical robustness making the system stable, repeatable, and highly accurate results at scale.

Accordingly, the project’s core hypotheses are considered validated, with the recommendation that future deployment phases prioritise mechanical hardening, environmental protection, and structured operational protocols to fully realise long-term industrial reliability and scalability.

4.6 Deliverables

All contracted milestones were achieved.

While certain subsystems (particularly width and rebar sensors) experienced shorter validated pilot windows due to physical disturbances caused by sleeper contact with sensor mounts, the milestones themselves were completed. The shortened validation windows were the result of mechanical impacts rather than technical infeasibility, and performance during stable calibration periods met the intended objectives.

Importantly, the system has been operating continuously in the plant environment since November 2025 without interrupting production, exceeding the originally defined pilot expectations.

4.6.1 Contracted Deliverables

The contracted deliverables and their status are summarised below:

Milestone 1 – Establish KPIs and Develop Defect Dictionary
Delivered:

- Report detailing baseline KPIs and evaluation metrics.
- Defined defect dictionary covering defect characteristics and scope selection.

Milestone 2 – Sensor Selection and ML/DL Model Research
Delivered:

- Report outlining hardware and software selection rationale.

Milestone 3 – Lab-Based Prototype IIoT System
Delivered:

- Functional prototype developed and tested at Swinburne.
- System presentation demonstrating operation.

Milestone 4 – AI-Based Defect Detection Development
Delivered:

- Report summarising AI model development, testing, and evaluation results.
- Integrated AI detection within IIoT architecture.

Milestone 5 – Pilot Testing and Final Report
Delivered:

- Pilot testing report including performance evaluation and improvement recommendations.
- Fully deployed prototype IIoT system (hardware and software).
- Handover of source code, trained models, datasets, documentation, and system operation guidelines.

All contracted deliverables were completed.

4.6.2 Modified Deliverables (Pivots)

There were no formal scope reductions. However, certain practical adjustments occurred during deployment:

- Due to mechanical impacts from sleeper handling, some sensors required reinforcement and recalibration during the pilot phase.
- As a result, the formal validation window for certain subsystems (particularly width and rebar detection) was shorter than originally anticipated.

These adjustments were operational rather than conceptual and did not alter the project's objectives or final deliverables.

4.6.3 Additional Deliverables (Not Originally Contracted)

The project delivered several additional capabilities beyond the original contract scope:

1. **Sleeper Pattern Detection via AI Vision**
   Automated classification of sleeper surface patterns (e.g., Smooth, Woodgrain, Stackstone, Block-faced) with 99.75% accuracy.
2. **Sleeper Colour Detection via AI Vision**
   Automated colour classification (Plain, Charcoal, Jasper, Yellow) with 94.23% accuracy.
3. **Graphical User Interface (GUI)**
   The original concept included installation of a signal tower and alarm to indicate defect detection.
   At the client's request, this was expanded to include the development of a full Graphical User Interface enabling:
   a. Real-time defect visualisation.
   b. Display of dimensional measurements.
   c. Review of inspection results.
   d. Historical traceability.

The GUI underwent a second design iteration after the client requested visibility of **all system-generated values**, not only defect events. This significantly enhanced transparency, traceability, and operator engagement.

4. **Sleeper Length Estimation via Transit Time Measurement**
   The system was enhanced to measure the time taken for each sleeper to pass through the inspection frame, enabling estimation of sleeper length. This feature was not part of the original contracted deliverables.

These additions increased the functional value of the system beyond the initial scope and strengthened its integration into plant operations.

4.7 Intellectual Property (IP)

4.7.1 Planned Project IP

The Project Agreement defined certain project outputs as Project IP, including the prototype inspection system design and the defect detection algorithm developed during the project.

While the project did not aim to produce patentable inventions, several forms of intellectual property were generated through the development and deployment of the automated inspection system. These primarily take the form of software, system design, datasets, and technical documentation developed specifically for the pilot implementation.

Key Project IP generated through the project includes:

Prototype Automated Inspection System Design
The integrated system architecture developed for automated quality control, including the configuration of industrial cameras, lighting, dimensional sensing, and computing infrastructure within a production inspection frame. The design specifies the physical layout, sensor positioning, data acquisition workflow, and integration with the inspection software environment.

Defect Detection Algorithm and AI Model Configuration
The machine vision pipeline and model configuration used to detect visual defects. This includes model architecture configuration, training workflow, preprocessing methods, inference logic, and associated software modules used to perform automated defect detection in a live production environment.

Inspection Software and Graphical User Interface (GUI)
Custom software developed to operate the inspection system, visualise inspection outputs, and present defect detection results to operators. This includes software modules for image acquisition, data processing, visualisation, and system control.

Training and Evaluation Datasets
Curated image datasets collected during the pilot deployment and used to train, validate, and evaluate the AI defect detection models. These datasets include labelled examples of production units and associated defect annotations used to support model development.

System Integration and Configuration Documentation
Technical documentation developed during the project describing system installation, calibration procedures, operational workflows, and maintenance processes required for ongoing system operation.

These items constitute the core Project IP associated with the development of the automated inspection capability.

4.7.2 Background IP

Certain knowledge inputs used during the project constitute Background IP owned by Sunset Concrete.

This includes the **defect classification framework (“defect dictionary”)** used to define relevant defect types and inspection criteria. This domain knowledge informed the development of the AI training dataset and model labelling framework but was not created as part of the project.

Ownership of this Background IP remains with Sunset Concrete.

4.7.3 Changes to Planned Project IP

There were no changes to the planned IP position during the project. The project generated the expected system design, software components, and supporting datasets as described in the Project Agreement.

4.7.4 Protection and Treatment of Project IP

The protection of Project IP generated through this project is outside the scope of this project and is within Sunsets management. Below are some suggestions to manage and protect IP :

- copyright protection for software code and documentation
- confidential information protections for algorithms, datasets, and system configurations
- contractual ownership and usage rights defined within the Project Agreement.
- a patent led by Sunset where applicable

At project completion, the relevant Project IP components — including software code, trained models, datasets, and system documentation — were formally transferred to Sunset Concrete as part of the agreed project deliverables.

These assets enable Sunset Concrete to operate, maintain, and further develop the automated inspection capability within its manufacturing operations.

Future exploitation of the Project IP may include continued operational use within production environments, further refinement of AI models using additional production data, and potential adaptation of the system for other precast manufacturing applications.

Part E - Next Steps – Commercialisation & Impact Plan

This section outlines the strategy to drive awareness, adoption, and long-term industry impact of the IIoT-enabled automated quality control system developed and piloted at Sunset Concrete.

5.1 Driving Awareness, Adoption and Application

To promote awareness and uptake within the precast and broader construction manufacturing sector, the following dissemination activities are planned:

- **Industry Article – Precast Association / Inside Construction**
  An article was published in the next edition of *Inside Construction* (Precast Association) outlining:
    - The industry challenge of manual QC limitations;
    - The design and pilot deployment of the IIoT system;
- **Community of Practice Presentation (17 February 2026)**
  A technical session titled:
  "AI and Industrial IoT for Smarter Quality Control of Concrete Sleepers"
    - System architecture
    - Demystifying AI
    - Real production data
- **Peer-Reviewed Publications (2 in preparation)**
  Two academic journal papers are currently being prepared, focusing on:
    - AI-based defect detection in live precast manufacturing environments.
    - Integration of IIoT sensing, machine vision, and industrial workflows for automated QC.

These publications will strengthen credibility, support future funding pathways, and position the project as a reference case for industrial AI adoption.

5.2 Ongoing Accessibility and Maintenance

Knowledge Transfer

A structured knowledge transfer process will be completed with Sunset Concrete, including:

- Technical walkthrough of system architecture;
- Training on operation, calibration, and troubleshooting;
- Documentation for maintenance and reconfiguration;
- Transfer of source code, AI models, and datasets.

Sunset Concrete is responsible for ongoing system operation and maintenance following formal handover.

Ongoing Maintenance Requirements

The system requires:

- Periodic mechanical inspection;
- Routine cleaning of camera and laser assemblies;
- Monitoring of AI model performance for rare defect classes;
- Basic IT infrastructure support for data storage and interface hosting.

No complex manufacturing or distribution activities are required at this stage, as the system is site-installed. Future scaling would require hardware replication and configuration support.

5.3 Assets Required for Industry Access

The assets enabling adoption include:

Capital Assets

- Industrial cameras and lighting assemblies;
- Laser-based dimensional sensors;
- Rebar detection sensors;
- Inspection frame and mounting hardware;
- Computing hardware.

All hardware assets have been transferred to Sunset Concrete.

Digital Assets

- Source code (vision, dimensional sensing, rebar detection);
- Trained AI vision models;
- Dataset used for model training and evaluation;
- Graphical User Interface (GUI);
- System documentation and operating manuals.

These digital assets have also been transferred to Sunset Concrete.

Skills and Organisational Capabilities Required

For continued operation and broader industry adoption, the following competencies are required:

- AI and machine vision engineering;
- Industrial automation and sensor integration expertise;
- IT infrastructure and data management capability;
- Production engineering oversight;
- Change management and operator training capability.

Future commercial scaling would likely require a cross-functional team combining engineering, software development, installation technicians, and industry sales functions.

5.4 Handover of Outputs

At Project Close

Provided to Sunset Concrete:

- Source code and trained AI models;
- Datasets and structured logs;
- GUI software;
- System documentation (installation, calibration, maintenance);
- Physical hardware system.

All outputs were prepared and verified prior to handover.

Longer Term

For broader dissemination:

- Industry-facing summaries will be published;
- Peer-reviewed papers will make methodological insights accessible;
- Knowledge shared through community forums.

5.5 Hosting and Responsibility

Internal Hosting

Sunset Concrete is responsible for:

- Hosting the system on-site;
- Maintaining hardware and computing infrastructure;
- Managing operational data storage.

Broader Access

Academic outputs will be made available through publication channels. Industry summaries will be disseminated through professional bodies and industry networks.

5.6 Project Data, Materials and Equipment

After Project Close

- Hardware remains installed at Sunset Concrete.
- Digital assets have been formally transferred.

- Research data remains archived for academic analysis and publication purposes, subject to agreed confidentiality terms.

Longer Term

- Equipment remains operational within the production environment.
- Datasets may support future model refinement and research initiatives.

5.7 Follow-up Research Opportunities

Immediate Research Opportunities

- Dataset expansion for rare defect classes (e.g., cracks) to improve precision.
- Mechanical optimisation of width measurement stability.
- In-line rebar cover sensor

Longer-Term Research Directions

- Fully autonomous closed-loop quality control systems;
- Predictive defect analytics using longitudinal production data;
- Integration with ERP and MES systems for real-time production optimisation;
- Multi-factory federated learning for AI generalisation across plants;
- Expansion of the framework to other precast or concrete products.

5.8 Future Deployment at Sunset Concrete

Sunset Concrete has indicated its intention to install and operate the automated quality control system as part of the commissioning of its new production facility in Pakenham, Victoria. The relocation from the current Officer site provides an opportunity to integrate the inspection capability within the updated production layout and quality assurance processes.

Within the new facility design, Sunset Concrete has allocated space for the inspection system in the post-demoulding stage of the production workflow. This location enables sleepers to be inspected immediately after demoulding, allowing visual and dimensional defects to be identified before products proceed further through the handling, storage, and dispatch stages.

The company views the automated inspection capability as a strategic quality assurance tool supporting its broader objectives to strengthen product consistency, improve traceability of quality outcomes, and enhance customer confidence. Integration of automated inspection within routine operations is expected to support faster and more objective quality verification compared with manual inspection processes.

Following the project handover, Sunset Concrete will evaluate operational integration of the system within the new factory environment, including workflow alignment, operator procedures, and ongoing system maintenance. Lessons learned during the pilot deployment will inform refinement of installation, calibration, and inspection procedures.

The planned installation at the new facility represents an important step in transitioning the system from a pilot demonstration to a sustained operational capability within a live manufacturing environment.

Part F - Education and Training

N/A

Part G - Conclusions, Implications and Recommendations

7.1 Challenges Overcome and Remaining

One of the primary challenges encountered during the project was mechanical instability within the production environment. Sensor mounts were occasionally displaced due to unintended contact during sleeper handling, and a small vertical step between the demoulding area and the measurement zone introduced dynamic disturbances that affected width measurements. These issues required reinforcement of sensor assemblies, recalibration, and refinement of installation practices. Despite these challenges, the system continued operating without interrupting plant production, and stable validation datasets were successfully obtained—particularly for height measurement and rebar detection during calibrated periods.

Another challenge involved dataset imbalance for rare defect types (e.g., cracks), which affected precision metrics. This was mitigated through conservative detection strategies prioritising recall to avoid missed critical defects.

Remaining challenges relate primarily to mechanical hardening of the width measurement station, further expansion of rare defect datasets to improve AI precision, and formalisation of operational reporting protocols when hardware disturbances occur. These are engineering and deployment refinements rather than fundamental technical barriers.

7.2 What Went Well – Highlights

Several aspects of the project exceeded expectations:

- The system has operated continuously in a live industrial environment since November 2025 without disrupting production.
- The vision module achieved 99.75% accuracy for pattern classification and 94.23% for colour classification under real factory conditions.
- Defect detection achieved high recall (≥0.93) for major defect classes and 100% recall for cracks.
- Height measurement demonstrated exceptional stability, with 98.8% of 12,898 sleepers falling within ±2 mm tolerance under production conditions.
- Rebar detection achieved a very low error rate (0.09%) during stable calibration periods.
- Additional functionalities—pattern recognition, colour classification, GUI development, and length estimation—were delivered beyond the original contracted scope.

A major highlight is the successful integration of AI, dimensional sensing, and IIoT logging into a single operational system deployed in a working precast manufacturing plant.

7.3 Key Learnings for Future Research and Commercialisation

Several important learnings emerged:

1. **Mechanical robustness is as critical as algorithmic performance.**
   Industrial AI systems must be physically protected and engineered for real plant interaction. Mechanical design should be prioritised alongside model development.

2. **High recall is more important than high precision in safety-critical defect detection.** From a quality assurance perspective, avoiding missed defects outweighs the operational cost of occasional false positives.

3. **Operator engagement and communication protocols are essential.** Successful deployment requires clear reporting processes for sensor disturbances and collaborative integration with plant workflows.

4. **Real production data is invaluable.** Continuous data capture under live conditions enabled meaningful validation and revealed practical constraints that laboratory testing alone would not expose.

5. **Scope naturally expands when value becomes visible.** The addition of GUI visualisation, colour and pattern detection, and length estimation demonstrates that once digital infrastructure is established, incremental capabilities can be added with relatively low marginal cost.

These learnings provide a strong foundation for future research projects and commercial deployments, particularly in scaling the system across additional manufacturing sites and extending AI-driven quality control to broader precast and construction applications.

Part H- Publication List

- **Industry Article – Precast Association / Inside Construction**
  An article was published in the next edition of *Inside Construction* (Precast Association) outlining:
    - The industry challenge of manual QC limitations;
    - The design and pilot deployment of the IIoT system;
- **Peer-Reviewed Publications (2 in preparation)**
  Two academic journal papers are currently being prepared, focusing on:
    - AI-based defect detection in live precast manufacturing environments.
    - Integration of IIoT sensing, machine vision, and industrial workflows for automated QC.

Part I - Appendices and Attachments

Proof of concept system deployed in the factory

Graphical User Interface Version 1

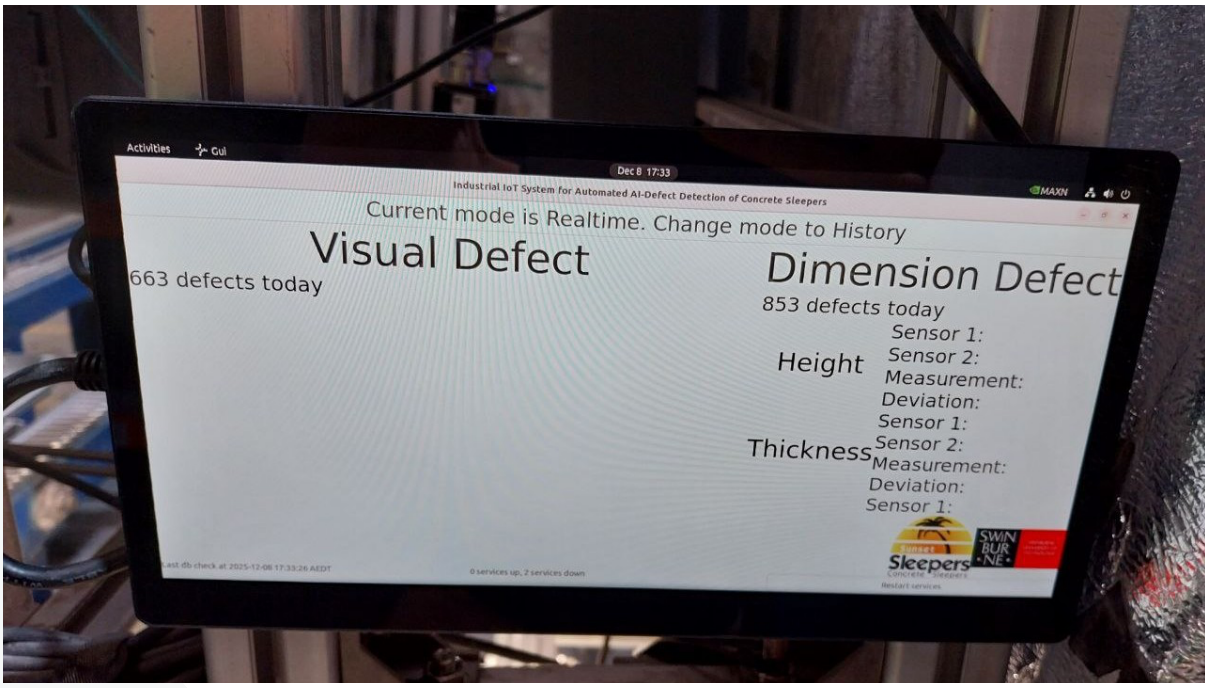

Activities Gui

Dec 8 17:33

Industrial IoT System for Automated AI-Defect Detection of Concrete Sleepers

Current mode is Realtime. Change mode to History

Visual Defect

663 defects today

Dimension Defect

853 defects today

Height
Sensor 1:
Sensor 2:
Measurement:
Deviation:

Thickness
Sensor 1:
Sensor 2:
Measurement:
Deviation:

Sensor 1:

Last db check at 2025-12-08 17:33:26 AEDT

0 services up, 2 services down

Restart services

Modified Graphical User Interface Version 2 based on Feedback from Sunset

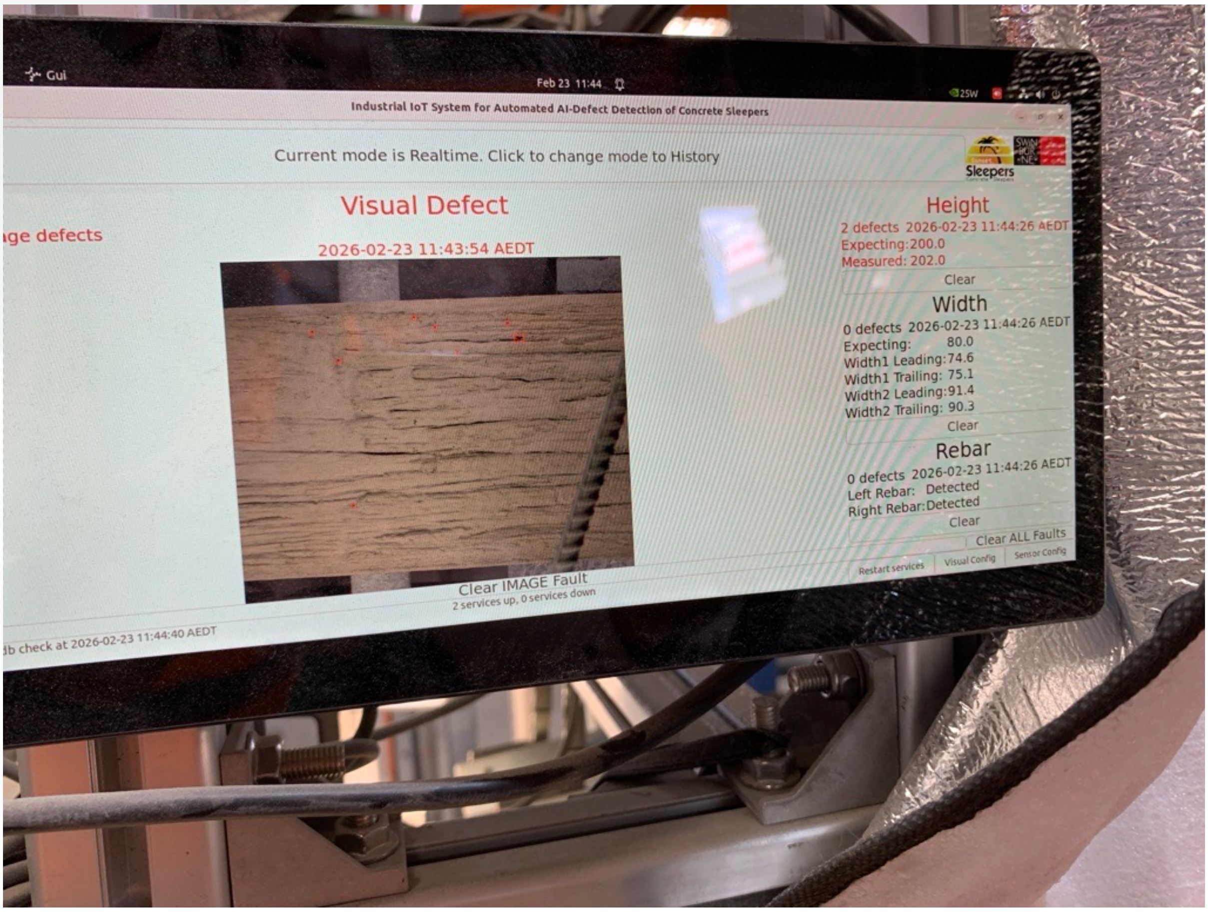

Current mode is Realtime. Click to change mode to History

ge defects

Visual Defect

2026-02-23 11:43:54 AEDT

Height

2 defects 2026-02-23 11:44:26 AEDT
Expecting:200.0
Measured: 202.0

Width

0 defects 2026-02-23 11:44:26 AEDT
Expecting: 80.0
Width1 Leading:74.6
Width1 Trailing: 75.1
Width2 Leading:91.4
Width2 Trailing: 90.3

Rebar

0 defects 2026-02-23 11:44:26 AEDT
Left Rebar: Detected
Right Rebar:Detected

2 services up, 0 services down

db check at 2026-02-23 11:44:40 AEDT

Video link that shows the problematic with the sleeper bouncing [Link Expire on 28th of May 2026]

https://liveswinburneeduau-my.sharepoint.com/:v:/g/personal/fmarti_swin_edu_au/IQCqHQWSDXAcTIhaIOWu-Z5sAR4ynjaZ_b_ABcL6AkjhGaM?nav=eyJyZWZlcnJhbEluZm8iOnsicmVmZXJyYWxBcHAiOiJPbmVEcml2ZUZvckJ1c2luZXNzIiwicmVmZXJyYWxBcHBQbGF0Zm9ybSI6IldlYiIsInJlZmVycmFsTW9kZSI6InZpZXciLCJyZWZlcnJhbFZpZXciOiJNeUZpbGVzTGlua0NvcHkifX0&e=aN5lEE