

Article

Hardware Integrity Verification Through Forensic Reasoning Using Manufacturing Images

Eyal Weiss 

Technology Department, Cybord.ai, Tel Aviv 6744332, Israel; eyal.w@cybord.ai; Tel.: +972-50-6292458

Abstract

Modern electronic assembly relies on complex global supply chains, making verification of component authenticity and consistency with the intended design increasingly important. Existing automated inspection approaches typically formulate this as classification or anomaly detection, providing limited insight into the physical evidence behind their conclusions and often failing to distinguish expected manufacturing variation from hardware integrity events. This paper presents a scenario-based method that formulates component verification as an evidence-based reasoning process. Semantic observations and learned visual evidence extracted from standard manufacturing images are evaluated against expected observations for candidate scenarios, including normal production evolution, approved AVL substitutions, unexpected component changes, and counterfeit-related events. The method was developed using more than 6.5 billion component images from high-volume SMT manufacturing. Quantitative evaluation across five controlled dual-population datasets comprising 1001 samples achieved 97.45% evidence-discrimination accuracy for manufacturing changes involving MPN, manufacturer, and production site. In a separate ground-truthed experiment, 19 of 20 counterfeit components were correctly detected within 200 components, with one false positive, corresponding to 99.0% accuracy. Longitudinal analysis of 429.6 million inspected components across 15 active EMS sites showed distinct behavior of marking-content changes: marking-to-MPN mismatches decreased by approximately 95% from 2023 to 2026 YTD, whereas date-code mismatches remained comparatively stable, illustrating how different changes within the same marking support different manufacturing scenarios and operational responses. The results demonstrate how visual and semantic evidence can jointly distinguish legitimate manufacturing variation, component substitutions requiring verification, and counterfeit-related inconsistencies. The methodology provides a scalable basis for component-level hardware assurance while creating a traceable digital record of the physical evidence supporting each assessment.



Academic Editors: Ferdinando
Di Martino and John A. Kalomiros

Received: 3 August 2026

Revised: 14 September 2026

Accepted: 22 September 2026

Published: 24 September 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: hardware integrity verification; hardware assurance; electronic assembly manufacturing; forensic reasoning; component authentication; explainable inspection; visual evidence; SMT inspection

1. Introduction

Electronic assemblies increasingly rely on complex global supply chains involving multiple component manufacturers, distributors, contract manufacturers, and logistics providers [1,2]. As a result, verifying that every assembled component is authentic, approved, and consistent with the intended design has become an essential aspect of electronic

assembly manufacturing [3–6]. Hardware integrity violations may arise from counterfeit components, unauthorized substitutions, remarked devices, recycled parts, supply chain mix-ups, manufacturing errors, storage conditions, or other events that cause the assembled hardware to differ from its intended manufacturing specification [5,7–13]. Figure 1 summarizes the proposed concept. Similar visual evidence may support different conclusions depending on the expected manufacturing scenario. Instead of searching for anomalies alone, the framework evaluates which hardware integrity scenario is most consistent with the observed evidence.




Figure 1. Scenario-based interpretation of manufacturing images. Rather than detecting anomalies alone, the proposed methodology determines which manufacturing or hardware integrity scenario best explains the observed visual evidence, enabling component-level hardware assurance.

Current approaches to hardware integrity verification combine documentation, supply chain traceability, electrical testing, destructive analysis, and forensic inspection. International standards, including SAE AS6171 and SAE AS5553 [6,14–16] for counterfeit detection, define structured inspection procedures for evaluating suspect components. These procedures are fundamentally hypothesis driven. Rather than prescribing a single inspection method, they collect different forms of evidence depending on the suspected scenario [6,14–16]. For example, resurfacing may be investigated using solvent testing, die replacement through decapsulation, and unauthorized substitutions through comparison with authentic reference devices. The objective is not simply to classify a component as authentic or counterfeit, but to determine which explanation is most consistent with the available evidence.

Although these inspection techniques differ considerably, much of the required evidence originates from the external appearance of the component. Manufacturing processes produce characteristic signatures in package geometry, body morphology, lead finish, logos, marking implementation, font characteristics, Pin-1 indicators, and other physical attributes. Likewise, resurfacing, remarking, refurbishment, unauthorized replacement, and environmental degradation frequently alter one or more of these visual signatures. Consequently, manufacturing images already contain a substantial portion of the information required for hardware integrity assessment [17–20].

Recent advances in computer vision and deep learning have substantially expanded the capabilities of automated visual inspection in electronics manufacturing. Deep learning methods have been applied to PCB and PCBA inspection for component recognition, defect detection, and hardware assurance, while anomaly-detection and self-supervised approaches have increasingly been used to identify deviations from learned normal manufacturing populations [10,21–25]. These approaches demonstrate the ability of learned visual representations to detect subtle physical differences under industrial imaging condi-

tions. However, most existing methods formulate the inspection task primarily as object or defect classification, detection, or deviation from a learned normal distribution [7,9,10,17,22–28]. Such formulations identify what is different, but do not inherently determine why it is different or whether the observed difference represents normal manufacturing evolution, an approved supply chain change, a process error, or a hardware integrity violation. This distinction motivates the scenario-based approach developed in this work, in which multiple independent observations are interpreted jointly as evidence for competing manufacturing and hardware integrity scenarios.

This work is based on the hypothesis that hardware integrity verification can be formulated as an evidence-based reasoning problem operating on manufacturing images. Rather than attempting to directly classify components as authentic or counterfeit, the proposed framework extracts multiple independent pieces of forensic visual evidence, each describing a different physical characteristic of the component. These observations are then evaluated against the expectations associated with candidate hardware integrity scenarios, in a manner analogous to forensic inspection. Figure 2 illustrates the relationship between conventional forensic inspection and the proposed methodology.




Figure 2. Comparison between conventional forensic inspection and the proposed method. Both approaches evaluate hardware integrity by collecting multiple independent observations and assessing their consistency with candidate hardware integrity scenarios. The proposed method automates evidence generation on all components on the board while preserving and automating the forensic reasoning process.

The proposed method consists of three principal stages. First, independent virtual sensors extract semantic observations and learned visual evidence from standard manufacturing images [17]. Second, the extracted evidence is evaluated against the expected observations associated with candidate hardware integrity scenarios, including normal manufacturing evolution, approved supply chain changes, manufacturing process events, and hardware integrity violations. Finally, the accumulated evidence is integrated into an explainable hardware integrity assessment that identifies both the most consistent scenario and the supporting physical observations.

The principal contribution of this work is therefore not a new binary counterfeit classifier, but an evidence-integration framework that determines which manufacturing or hardware integrity scenario best explains the observed component. This distinction is evaluated experimentally using controlled changes in manufacturer, MPN, and production site, together with an independently ground-truthed counterfeit experiment, allowing both legitimate manufacturing changes and hardware integrity violations to be examined within the same reasoning framework.

2. Scenario-Based Hardware Integrity Assessment

Hardware integrity verification is fundamentally a process of evaluating competing explanations for the observed hardware [3,6,14–16]. Rather than determining whether a component is simply authentic or counterfeit, the objective is to identify the scenario that best explains its observed physical characteristics. While counterfeit detection is an important application, most components inspected during production are expected to be correct. The verification process must therefore distinguish between normal manufacturing variation, expected supply chain changes, manufacturing process events, and genuine hardware integrity violations.

Each hardware integrity scenario defines a characteristic set of expected observations. Components originating from the same reel are expected to exhibit nearly identical package geometry, package body characteristics, manufacturer signatures, marking style, Pin-1 implementation, lead geometry, and marking content [17]. In contrast, the introduction of a new manufacturing lot is expected to preserve most physical characteristics while changing only specific marking fields, such as the lot number or date code. Similarly, an approved AVL substitution introduces manufacturer-specific differences while maintaining compatibility with the approved package and manufacturing specification [1,29–32]. Hardware integrity violations, including unauthorized substitutions, resurfaced packages, remarked devices, refurbished components, and counterfeit parts, produce different combinations of expected observations reflecting their underlying physical modifications.

Unlike conventional classification approaches, the proposed method evaluates whether the collected evidence is consistent with the expectations associated with each candidate scenario. Every evidence item therefore supports, contradicts, or remains neutral with respect to a given hypothesis. Hardware integrity assessment is obtained by combining these independent observations to identify the scenario that best explains the observed component.

Representative scenarios and their expected observations are summarized in Table 1. Individual observations are rarely sufficient to confirm or reject a scenario. Instead, the assessment relies on the combined consistency of multiple independent sources of visual evidence.

Table 1. Representative hardware integrity scenarios and their expected visual observations.

Scenario	Expected to Remain Consistent	Expected to Change
Same reel	Package, body fingerprint, manufacturer fingerprint, logo, marking style, font, pin-1, leads, surface texture, lot code, date code	None (except minor manufacturing variation)
New manufacturing lot	Package, body fingerprint, manufacturer fingerprint, logo, marking style, font, pin-1, leads	Lot code, possibly date code
New date code	All physical characteristics	Date code only
Approved AVL substitution	Package compatibility, functional characteristics	Manufacturer logo, manufacturer fingerprint, body morphology, marking implementation
Package revision	Functionality, manufacturer identity	Package geometry, body morphology, pin configuration
Mixed reel	Multiple internally consistent populations	Unexpected coexistence of different visual populations within the same reel

Table 1. *Cont.*

Scenario	Expected to Remain Consistent	Expected to Change
Incorrect feeder/wrong component	Placement location	Package, markings, manufacturer, geometry inconsistent with expectation
Remarked component	Package geometry, manufacturer body characteristics	Marking text, font, alignment, engraving or printing characteristics
Resurfaced/blacktopped component	Internal package geometry	Surface texture, reflectance, marking quality
Refurbished component	Package identity	Lead condition, oxidation, scratches, solder residue
Counterfeit/unauthorized substitution	Varies depending on mechanism	Multiple inconsistencies across package, markings, manufacturer signatures, and production characteristics

3. Visual Evidence Extraction

Hardware integrity assessment requires observations describing different physical characteristics of each component. Rather than representing a component using a single feature vector or classification model, the proposed framework decomposes visual inspection into a collection of independent evidence extraction modules. Each module analyzes a specific physical characteristic and produces an observation that is subsequently evaluated against the expectations associated with candidate hardware integrity scenarios.

The evidence extraction models were developed using a dataset of more than 6.5 billion component images collected from Pick-and-Place and Automated Optical Inspection (AOI) systems operating on high-volume electronic assembly production lines. The images were grouped and labeled according to the physical characteristics represented by each evidence category, enabling each model to learn the normal variability associated with a specific visual attribute rather than complete component identities. To ensure deployment across different manufacturing environments, the models were trained to be robust to variations introduced by different machine vendors, imaging systems, optical configurations, and production conditions [17].

This modular organization follows the methodology used during forensic inspection, where different physical characteristics are examined independently before reaching an overall conclusion. Separating the evidence sources also improves interpretability, allowing each observation to support or contradict a scenario independently rather than contributing to an opaque classification score.

The method operates on standard manufacturing images acquired during electronic assembly production, including images from Pick-and-Place systems [20], Automated Optical Inspection (AOI) [17], incoming inspection [13], or similar imaging equipment. Following image normalization, individual evidence modules analyze different physical characteristics of the component. The principal categories of evidence are summarized in Table 2.

Each evidence module operates independently and produces an observation that remains interpretable throughout the reasoning process. New evidence sources can therefore be incorporated without modifying the overall scenario evaluation framework, making the methodology readily extensible as additional inspection capabilities become available.

Table 2. Visual evidence categories extracted from manufacturing images.

Evidence Category	Representative Observations	Typical Supported Scenarios
Package Geometry	Package dimensions, aspect ratio, pin count, package family	Package substitution, feeder errors, package revisions
Package Body Characteristics	Surface texture, molding signatures, cavity marks, body morphology	Manufacturer identification, counterfeit, package replacement
Manufacturer Signatures	Logos, proprietary markings, manufacturer-specific visual characteristics	AVL substitutions, unauthorized manufacturer changes
Marking Morphology	Font, layout, spacing, engraving or printing characteristics	Remarking, manufacturing variation
Marking Content	MPN, date code, lot code, country of origin	New lots, date-code changes, traceability
Pin-1 Characteristics	Location, geometry, manufacturing style, surface appearance	Package authentication, resurfacing, remarking
Lead Characteristics	Shape, finish, oxidation, scratches, deformation, solder residue	Refurbishment, degradation, reuse
Population Consistency	Similarity across components from the same production population	Same reel, mixed reels, counterfeit, unexpected variation

4. Scenario Evaluation

The extracted visual evidence is evaluated against the expectations associated with each candidate hardware integrity scenario. Since every scenario defines a different set of expected observations, the same evidence may support one scenario while contradicting another. For example, a change in the date code is expected when a new manufacturing lot is introduced but is inconsistent with the Same Reel scenario. Likewise, a different manufacturer signature is expected for an approved AVL substitution but would contradict both the Same Reel and New Manufacturing Lot scenarios.

The scenario evaluation therefore consists of comparing the complete set of extracted observations with the expected evidence pattern associated with each candidate scenario. Hardware integrity is determined by identifying the scenario whose expected observations are most consistent with the collected evidence. Because each evidence source remains independent throughout the evaluation, the resulting assessment is inherently explainable and identifies which observations support or contradict the final conclusion. Representative expectations for several common scenarios are summarized in Table 3.

Table 3. Representative expectations for several common scenarios.

Evidence	Same Reel	New Lot	MPN Change	Unauthorized
Package Geometry	Match	Match	May differ	May differ
Manufacturer Identity	Match	Match	May differ	May differ
Package Body Fingerprint	Match	Match	May differ	May differ
Marking Style	Match	Match	May differ	May differ
Marking Text	Match	Lot/Date Change	MPN Change	Match
Population Consistency	High	High	Unexpected population	May be inconsistent
Scenario Interpretation	Normal	Expected Change	Requires Verification	Concern

Scenario evaluation uses a predefined expectation table specifying the expected response of each evidence test for each candidate scenario. Individual evidence scores are compared with these expected patterns and combined into a scenario-level consistency score.

The scenario with the highest consistency is selected, while the individual evidence scores are retained to explain which observations support or contradict the resulting assessment.

The patterns shown are representative rather than exhaustive. Actual scenario evaluation combines the individual evidence scores described in Section 4. The scenario evaluation process is independent of the algorithms used to generate the individual evidence. Any image analysis, measurement, or future sensing technology capable of producing interpretable observations can be incorporated into the framework, provided that its outputs can be compared with the expectations associated with the evaluated scenarios.

5. Method

The proposed framework separates evidence generation from hardware integrity reasoning. Each evidence extraction module produces a standardized representation describing a single physical characteristic of the component. These representations are subsequently evaluated by the scenario-based reasoning framework described in Section 4.

Two complementary forms of evidence are employed throughout the framework. Physical attributes that can be directly interpreted, such as marking text, manufacturer identity, package family, date codes, and lot identifiers, are represented as semantic evidence. Visual characteristics that cannot be adequately described using predefined labels are represented using learned evidence embeddings.

5.1. Semantic Evidence

Semantic evidence represents structured information that can be interpreted directly without comparison. Examples include optical character recognition (OCR) results, manufacturer part numbers (MPNs), manufacturer identity, package family, date codes, lot identifiers, and country of origin.

These observations are evaluated directly against the expectations associated with each hardware integrity scenario. Depending on the evaluated scenario, identical values may be required, expected to change, or permitted to differ.

5.2. Learned Evidence Embeddings

Many categories of forensic evidence cannot be represented reliably using predefined measurements or discrete labels. Characteristics such as package body appearance, molding signatures, marking morphology, manufacturer fingerprints, Pin-1 implementation, lead appearance, and other manufacturing signatures are inherently visual. These characteristics are therefore represented using learned evidence embeddings.

Each evidence module transforms its corresponding image region into a compact numerical representation describing a single physical characteristic. Rather than representing an entire component using one feature vector, the proposed framework maintains an independent embedding space for each evidence category. This separation enables every embedding to specialize in a specific forensic observation while avoiding interference from unrelated visual information.

The embedding models are trained using billions of component images collected during normal SMT manufacturing. Instead of learning component identities or counterfeit classes, components are grouped according to the physical characteristic represented by each evidence category. Consequently, each embedding learns to preserve manufacturing consistency while separating meaningful physical differences relevant to its assigned observation. Because the training data originates from large-scale production environments, the learned representations naturally capture normal variation across manufacturing lines, suppliers, production lots, imaging systems, and package revisions.

The learned evidence modules are based on ResNet-type convolutional architectures, with the embedding representation adapted to the physical evidence being evaluated. For example, some modules operate on the complete component image, whereas others operate on localized regions such as the component marking. Rather than applying a fixed similarity threshold across all component types and evidence categories, similarity is evaluated relative to the natural variation observed within the corresponding reference population. The distribution of embedding distances within the reference population is characterized using robust statistics based on the median and deviation from the median, and candidate observations are evaluated according to their normalized distance from this population. Acceptance thresholds are defined as multiples of the observed population deviation; for example, a threshold may be selected several deviations from the population center depending on the required sensitivity. Consequently, the decision boundary adapts to the intrinsic visual variability of each reference population rather than relying on a globally fixed similarity score.

Hardware integrity assessment is not performed directly from the embeddings themselves. Instead, similarity measures quantify the consistency between corresponding evidence embeddings and provide observations that are subsequently evaluated by the scenario reasoning framework. This separation between evidence generation and evidence interpretation enables the same embedding representations to support multiple hardware integrity scenarios without modification.

6. Experimental Validation

The proposed framework was evaluated using production images acquired during standard SMT manufacturing. The dataset consists of billions of component images collected in-line from Pick-and-Place and Automated Optical Inspection (AOI) systems across multiple manufacturing sites, customers, component manufacturers, package families, and production lots. Rather than relying on curated benchmark datasets, the models were developed using large-scale production imagery, while the validation described below was performed using controlled manufacturing populations with known changes and an independently ground-truthed counterfeit experiment.

The evaluation combines three complementary forms of analysis. First, controlled manufacturing populations with known changes in MPN, manufacturer, and production site are used to quantify the discrimination capability of the underlying evidence layer. Second, an independently ground-truthed counterfeit experiment provides direct quantitative validation of one hardware integrity scenario. Third, longitudinal production data are used to examine how different marking-content evidence patterns and their associated scenario interpretations evolve during sustained manufacturing operation. Together, these analyses evaluate evidence discrimination, scenario-level decision making, and the behavior of the framework under large-scale production conditions.

6.1. Quantitative and Scenario-Based Validation

To quantitatively evaluate the evidence used by the scenario-based reasoning framework, two complementary validation experiments were performed. The first consisted of five controlled dual-population datasets comprising 1001 component samples and representing known manufacturing changes rather than arbitrary visual perturbations. The evaluated cases included changes in both MPN and manufacturer, a change in MPN within the same manufacturer, and a production-site change while preserving both manufacturer and MPN. These cases exercise different combinations of visual and semantic evidence used by the framework to distinguish normal manufacturing evolution, component substitutions, and other manufacturing changes.

A total of 1001 component samples were evaluated across the five controlled dual-population datasets. Each sample was evaluated both against its own reference population and against the alternative population representing the known manufacturing change. All 1001 samples were correctly matched to their own reference population, corresponding to a 100% within-population acceptance rate with no false rejections. When evaluated against the alternative population, 950 samples (94.9%) were correctly rejected, while 51 samples (5.1%) were incorrectly accepted. Considering both within-population acceptance and cross-population rejection, the resulting evidence-discrimination accuracy was 97.45%, with 100% precision and 94.9% recall when rejection of the alternative population is treated as the positive class.

A second experiment evaluated a true counterfeit scenario using components with independently established ground truth. A reference population of 100 authentic components was used. The evaluation set contained 200 components, comprising 180 authentic components and 20 counterfeit components obtained from an external laboratory and mixed into the authentic population. The system correctly identified 19 of the 20 counterfeit components and correctly accepted 179 of the 180 authentic components, corresponding to 99.0% classification accuracy, 95.0% counterfeit precision, 95.0% counterfeit recall, and 99.44% specificity.

Table 4 illustrates that a visual population difference is not itself interpreted as evidence of counterfeiting. Instead, the scenario is determined from the consistency between visual observations, decoded marking information, and the expected manufacturing context. In Datasets 2, 3, and 16, the visual differences are accompanied by a decoded change in both manufacturer and MPN. The components therefore do not attempt to represent the expected component identity; the resulting scenario is a manufacturer/MPN substitution, which can subsequently be checked against the applicable Approved Vendor List (AVL) to distinguish an approved change from an unauthorized substitution.

Table 4. Evidence-based interpretation of the manufacturing-change and counterfeit validation scenarios.

Dataset #	Ground-Truth Scenario	Visual Evidence	Semantic Evidence	Scenario Interpretation
2, 3, 16	Manufacturer and MPN change	All visual tests indicate a different population, including manufacturer-related evidence	Marking decoding identifies a different manufacturer and MPN	Manufacturer/MPN substitution; AVL determines whether the change is approved or unauthorized
13	Production-site change; same manufacturer and MPN	Multiple visual tests indicate substantial population differences	Manufacturer and MPN match; site code changes from M to K	Legitimate production-site change; visual differences are explained rather than interpreted as counterfeit
15	MPN change; same manufacturer	Component appearance and marking style remain consistent, while marking content changes	Manufacturer matches; decoded marking identifies a different MPN	MPN substitution; AVL determines whether the new MPN is approved
Counterfeit	Blacktopped and remarked counterfeit	Package appearance and marking style differ; contact geometry shows smaller differences	Claimed component identity and general marking format remain consistent with the expected component	Physical evidence contradicts the claimed identity; counterfeit scenario

Datasets 13 and 15 demonstrate why semantic interpretation of the marking is particularly important. In Dataset 13, substantial visual differences occur despite an unchanged manufacturer and MPN. The marking parser identifies a production-site code change from “M” to “K”, providing a manufacturing explanation for the visual population shift and preventing it from being interpreted as counterfeit (Figure 3). Dataset 15 exhibits the complementary case: the manufacturer-related appearance and marking style remain consistent, but the marking content changes. Decoding the marking identifies a different

MPN from the same manufacturer, allowing an MPN substitution to be distinguished from a normal date-, lot-, or production-related marking change.




Figure 3. Same manufacturer and MPN produced at two different manufacturing sites.

Ground truth for the controlled datasets was established independently from the known component identity and manufacturing information associated with each population, including manufacturer, MPN, and production-site information where applicable. The two populations within each dataset were defined before similarity analysis, allowing both population discrimination and the resulting evidence pattern to be evaluated against the known manufacturing change. For the counterfeit experiment, ground truth was established independently from the authentic and physically counterfeit components supplied for the experiment.

6.2. Expected Manufacturing Evolution

The following cases provide detailed examples of how the evidence patterns summarized in Table 4 and the broader scenario framework are interpreted for representative manufacturing and hardware integrity events. Figure 4 presents a component originating from a new manufacturing lot. The framework identifies that only the manufacturing date code differs from the reference population, while all remaining evidence remains consistent. The observations therefore support the New Manufacturing Lot scenario and do not indicate a hardware integrity violation.

6.3. Approved Manufacturer Substitution

Figure 5 illustrates an approved manufacturer substitution. In this case, the visual evidence indicates that the component belongs to a different manufacturing population, with differences in package appearance, marking characteristics, and manufacturer-related features. Semantic decoding independently identifies a different manufacturer and MPN. These observations are therefore interpreted as a component substitution rather than as evidence of counterfeiting. The identified component is subsequently evaluated against the customer AVL; when the manufacturer and MPN correspond to an approved alternative, the change is classified as an authorized substitution.




Figure 4. Example of a component belonging to a new manufacturing lot. The framework identifies that only the manufacturing date code differs from the reference population, while all remaining evidence remains consistent. The observations therefore support the “New Manufacturing Lot” scenario rather than indicating a hardware integrity violation.

6.4. Counterfeit Component

The counterfeit case, shown in Figure 6, provides a different evidence pattern. The counterfeit components had been blacktopped and remarked to reproduce the expected component identity and general marking format. Consequently, the decoded manufacturer and MPN were consistent with the expected component and did not indicate a legitimate manufacturer or MPN substitution. However, the package appearance and marking style differed from the authentic reference population, while smaller differences were also observed in the lead/contact geometry, likely as a result of component reprocessing. The combination of a matching claimed identity with inconsistent physical manufacturing evidence is therefore consistent with a counterfeit scenario.

This reasoning pattern was evaluated quantitatively in the independently ground-truthed counterfeit experiment described in Section 6.1. Of the 20 counterfeit components, 19 were correctly identified, while 179 of 180 authentic components were correctly accepted, resulting in 99.0% overall classification accuracy. This result provides direct validation of the counterfeit scenario rather than only of the underlying population-discrimination capability.




Figure 5. Example of an approved AVL replacement. The manufacturer and MPN differ from the reference population, while the remaining evidence is consistent with an approved alternate component. The method recognizes the substitution as an expected supply chain event rather than a hardware integrity violation.

6.5. Unexpected Component Substitution

Figure 7 illustrates an unexpected MPN change within the same manufacturer. In this case, package appearance and manufacturer-related visual characteristics remain consistent with the reference population, while the marking content changes. Semantic decoding identifies that the manufacturer remains unchanged but the decoded MPN differs from the expected component. The evidence therefore indicates a component substitution rather than counterfeiting. The decoded MPN is subsequently evaluated against the expected BOM/AVL information to determine whether the substitution is authorized or represents an unexpected hardware change.

6.6. Computational Performance

The system is designed for production-scale operation using a pipelined processing architecture. In a representative deployment using a consumer-grade NVIDIA RTX 5090 GPU, a PCBA containing approximately 3000 components can be fully analyzed and compared with its reference population at a sustained throughput of approximately one board every 40 s. This processing includes the analysis of all components on the board and their population-based evidence comparison. The pipelined implementation enables continuous operation alongside the manufacturing process.




Figure 6. Representative counterfeit evaluation. The framework compares the unit under test with a reference population and independently evaluates multiple evidence categories. Although several characteristics remain consistent with the reference population, multiple independent mismatches indicate that the observed evidence is inconsistent with expected manufacturing variation, resulting in a counterfeit hardware integrity assessment.

6.7. Longitudinal Production Observations

In addition to the controlled validation experiments, the behavior of the framework was examined using a large and heterogeneous set of routine production data collected over an extended period. The longitudinal dataset comprised 429.6 million inspected components produced across 15 active EMS manufacturing sites, representing 14.42 million assembled boards and 1692 distinct board part numbers. The inspected population included 13,244 distinct component MPNs from 27 component manufacturers and production associated with five major OEM customers. Sites without component-level inspection records were excluded from the analysis. These data are distinct from the broader 6.5-billion-image corpus used to develop the underlying visual evidence models; here, the objective was to examine the occurrence and evolution of detected discrepancies under sustained real-world manufacturing conditions.

The analysis focused on two automatically detected discrepancy categories that originate from the marking-content evidence used by the scenario-based reasoning framework: marking-to-MPN mismatch and date-code mismatch. As described in Sections 4 and 5, the framework does not treat a change in component marking as a single generic anomaly. Instead, the decoded marking fields are compared with the reference population to determine what information has changed, and the resulting evidence is interpreted according to the corresponding manufacturing scenario. In the present analysis, a change in the decoded MPN provides evidence of a component-identity change and therefore supports

an MPN-substitution scenario requiring verification, whereas a change confined to the manufacturing date code is consistent with normal production evolution, such as a new manufacturing lot. These two evidence patterns therefore provide a useful longitudinal example of how different changes detected within the same component marking lead to different scenario interpretations. For each month, the number of detected discrepancies was normalized by the corresponding inspected component volume, allowing comparison across periods with different production volumes.




Figure 7. Scenario-based evaluation of an unexpected MPN change. The package, manufacturer identity, and manufacturing signatures remain consistent with the reference population, indicating that the component originates from the same manufacturer. However, the detected MPN differs from the expected device, suggesting a possible inventory mix-up, unauthorized substitution, or use of an alternate component requiring verification.

The longitudinal analysis was performed after the initial manufacturer-specific reference-population learning and stabilization period. During early deployment, reference populations were still being accumulated and adapted to the normal variability associated with individual component manufacturers and production environments, resulting in larger transient fluctuations. This initialization period was excluded so that the longitudinal analysis reflects steady-state operation rather than initial population-building effects.

As shown in Figure 8, the two discrepancy categories exhibit markedly different longitudinal behavior. Following stabilization, the average marking-to-MPN mismatch rate decreased from 7.9% in 2023 to 3.3% in 2024, 0.8% in 2025, and 0.4% in 2026 YTD, corresponding to an approximately 95% relative reduction from the 2023 level. In contrast, the date-code mismatch rate remained comparatively stable, averaging 1.9%, 1.2%, 1.7%, and 1.3% over the same periods, respectively. The monthly data are shown directly rather

than only as annual aggregates in order to preserve the underlying temporal variability and demonstrate that the observed difference represents a sustained longitudinal trend.




Figure 8. Longitudinal production behavior of marking-to-MPN and date-code discrepancies across a heterogeneous manufacturing population. Monthly mismatch rates are normalized by the corresponding inspected component volume. The analysis includes production from 15 active EMS sites and begins after the initial manufacturer-specific reference-population learning and stabilization period. Marking-to-MPN discrepancies exhibit a pronounced long-term decline, whereas date-code discrepancies remain comparatively stable.

The different trajectories are particularly relevant to the scenario-based formulation of the framework because both observations originate from changes detected in the component marking, yet they represent different manufacturing scenarios. An MPN change is interpreted as a component-identity change requiring verification against the expected BOM/AVL, whereas a date-code change can be consistent with expected manufacturing evolution and is retained primarily as traceability evidence. Their operational consequences are therefore different. An MPN mismatch can trigger investigation and corrective action involving material selection, supplier information, feeder configuration, or manufacturing processes. Repeated feedback and correction can consequently reduce recurrence of this discrepancy class.

Date-code differences, in contrast, are normally preserved as part of the component traceability record for subsequent recall, investigation, or provenance analysis, without a corresponding objective to eliminate legitimate date-code variation. Consequently, no comparable long-term reduction would necessarily be expected.

These longitudinal observations are not presented as an additional measurement of classification accuracy, since individual discrepancy events across the complete production population were not independently ground-truthed. Instead, they provide complementary production-scale evidence of the behavior of the proposed framework under sustained and heterogeneous manufacturing conditions.

7. Discussion

The proposed method reformulates hardware integrity verification as an evidence-based reasoning problem rather than a component classification task. Instead of attempting to determine whether a component belongs to a predefined category, the methodology evaluates whether the observed physical evidence is consistent with the expectations associated with a candidate hardware integrity scenario. This formulation more closely reflects

established forensic inspection practice, where conclusions are derived from the combined interpretation of multiple independent observations rather than from a single measurement.

The validation results demonstrate two complementary aspects of the proposed framework. The five controlled manufacturing-change datasets evaluate the ability of the underlying evidence layer to distinguish populations associated with known changes in MPN, manufacturer, and production site, achieving 97.45% overall evidence-discrimination accuracy. Because the causes of these population changes are independently known, the resulting visual and semantic evidence can also be evaluated against the corresponding scenario interpretation. The same visual differences may therefore lead to different conclusions depending on the accompanying semantic evidence: a production-site change can explain substantial visual variation, an MPN or manufacturer change identifies a component substitution requiring AVL verification, whereas physical inconsistencies that contradict an otherwise matching claimed identity can indicate counterfeiting. The separate counterfeit experiment provides direct quantitative validation of the latter scenario, achieving 99.0% classification accuracy on 200 components with independently established authentic and counterfeit ground truth.

The longitudinal production analysis provides a complementary production-scale perspective on the scenario-based framework. Although both MPN and date-code discrepancies originate from marking-content analysis, they represent different manufacturing scenarios and operational responses. Their distinctly different long-term trajectories are therefore consistent with the framework's interpretation of an MPN change as an actionable component-identity discrepancy and a date-code change as primarily traceability evidence. The breadth of the analyzed production population further reduces the likelihood that these trends are specific to a single component family, manufacturer, product, or manufacturing site. These results are interpreted as production-scale evidence of framework behavior and operational relevance rather than as an additional measure of scenario-classification accuracy.

This modular organization allows additional evidence sources to be incorporated without modifying the overall reasoning framework. Although this work focuses on manufacturing images, future evidence sources such as X-ray inspection, acoustic microscopy, electrical testing, material characterization, or supply chain documentation can be integrated naturally as additional evidence modules.

The method also extends hardware integrity verification beyond counterfeit detection. Most inspected components encountered during manufacturing represent expected production events, including new manufacturing lots, approved AVL substitutions, package revisions, or normal process variation. Evaluating these scenarios within the same reasoning framework enables expected manufacturing evolution to be distinguished from conditions requiring further investigation, reducing unnecessary false alarms while preserving explainability.

The proposed methodology is limited by the information available in the observed evidence and by the quality and representativeness of the reference populations. Variations in imaging conditions, limited reference populations, or unusual manufacturing variation may reduce the reliability of individual evidence tests and increase false indications. Similarly, sophisticated modifications that remain within the natural visual variability of the reference population may not be detected. The present quantitative evaluation covers several representative manufacturing-change scenarios, including MPN, manufacturer, and production-site changes, together with an independently ground-truthed counterfeit case. It does not, however, constitute exhaustive statistical validation of every scenario defined by the method.

The longitudinal production analysis characterizes detected discrepancy patterns and their evolution over time but does not provide independently established ground

truth for every individual production event. Broader validation across additional component families, manufacturers, production sites, and hardware integrity scenarios remains necessary. In cases of insufficient or conflicting evidence, the framework can support escalation to complementary inspection methods such as electrical testing, X-ray inspection, microscopy, or destructive analysis. The method should therefore be viewed as an extensible evidence-integration methodology rather than a replacement for established hardware-assurance practices.

8. Conclusions

This paper presented a scenario-based methodology for hardware integrity verification using visual evidence extracted from standard electronic assembly manufacturing images. Rather than treating hardware integrity as a binary classification problem, the proposed approach evaluates multiple independent observations against the expected evidence patterns associated with candidate manufacturing and hardware integrity scenarios. This enables expected production evolution, approved supply chain changes, and potential hardware integrity violations to be interpreted within a common and explainable framework.

The methodology combines directly interpretable semantic observations with learned visual evidence while maintaining a clear separation between evidence generation and scenario interpretation. Quantitative evaluation across five controlled dual-population datasets comprising 1001 samples demonstrated 97.45% evidence-discrimination accuracy for known changes involving MPN, manufacturer, and production site. The scenario interpretation of these changes demonstrates how semantic and visual evidence can distinguish legitimate manufacturing variation from component substitutions requiring verification. In a separate experiment using independently established ground truth, the framework correctly classified 198 of 200 components in a mixed population containing 20 physically counterfeit and 180 authentic components, corresponding to 99.0% counterfeit-classification accuracy.

Longitudinal analysis across 429.6 million inspected components further demonstrated production-scale behavior consistent with the scenario framework: marking-to-MPN and date-code changes exhibited distinctly different long-term trends corresponding to their different manufacturing interpretations and operational consequences.

Applied at manufacturing scale, this approach enables components assembled on PCBs to be automatically inspected, compared with their expected populations, and documented as part of the normal production process. The resulting digital record preserves both the component-level assessment and the physical evidence supporting it, providing a practical basis for scalable and traceable hardware assurance. By transforming standard manufacturing images into structured forensic evidence, the proposed framework extends automated inspection from detecting visual differences toward understanding whether those differences are expected, explainable, or indicative of a hardware integrity concern.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: Author was employed by Cybord.ai. The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Razak, G.M.; Hendry, L.C.; Stevenson, M. Supply chain traceability: A review of the benefits and its relationship with supply chain resilience. *Prod. Plan. Control* **2023**, *34*, 1114–1134. [\[CrossRef\]](#)
2. Pecht, M.; Dagorn, F.; Das, D. Managing Electronics Part Changes in the Supply Chain. *IEEE Trans. Compon. Packag. Manuf. Technol.* **2018**, *8*, 883–895. [\[CrossRef\]](#)
3. Tehranipoor, M.M.; Guin, U.; Forte, D. Counterfeit integrated circuits. In *Counterfeit Integrated Circuits*; Springer: Berlin/Heidelberg, Germany, 2015; pp. 15–36.
4. Gayialis, S.P.; Kechagias, E.P.; Papadopoulos, G.A.; Masouras, D. A Review and Classification Framework of Traceability Approaches for Identifying Product Supply Chain Counterfeiting. *Sustainability* **2022**, *14*, 6666. [\[CrossRef\]](#)
5. Richman, D.R.; Azarian, M.H.; Das, D.; Pecht, M. Analysis of standards-based counterfeit microelectronics detection methods. *IEEE Access* **2025**, *13*, 7691–7704. [\[CrossRef\]](#)
6. Azarian, M.H. An Overview of Risk-Based EEE Counterfeit Part Detection Based on SAE AS6171. In *ISTFA 2018: Proceedings from the 44th International Symposium for Testing and Failure Analysis*; ASM International: Almere, The Netherlands, 2018.
7. BMagaña, M.; Hernández-Urbe, Ó.; Cárdenas-Robledo, L.A.; Cantoral-Ceballos, J.A. Deep Learning Algorithms for Defect Detection on Electronic Assemblies: A Systematic Literature Review. *Mach. Learn. Knowl. Extr.* **2026**, *8*, 5. [\[CrossRef\]](#)
8. Shi, J.; Fang, X.; Liu, Z. Design of A Detection and Recognition System for Mounting Defects of Surface Mount Technology Components Based on Improved Convolutional Neural Networks. In *2021 IEEE International Conference on Artificial Intelligence, Robotics, and Communication (ICAIRC)*; IEEE: New York, NY, USA, 2021; pp. 16–19.
9. de Oliveira, G.G.; Vaz, G.C.; Andrade, M.A.; Iano, Y.; Ximenes, L.R.; Arthur, R. System for PCB Defect Detection Using Visual Computing and Deep Learning for Production Optimization. *IET Circuits Devices Syst.* **2023**, *2023*, 6681526. [\[CrossRef\]](#)
10. Bhattacharya, A.; Cloutier, S.G. End-to-end deep learning framework for printed circuit board manufacturing defect classification. *Sci. Rep.* **2022**, *12*, 12559. [\[CrossRef\]](#) [\[PubMed\]](#)
11. MTehranipoor, M.; Guin, U.; Forte, D. Physical tests for counterfeit detection. In *Counterfeit Integrated Circuits*; Springer: Berlin/Heidelberg, Germany, 2015; pp. 75–93.
12. Guin, U.; Tehranipoor, M. On selection of counterfeit IC detection methods. In *IEEE North Atlantic Test Workshop (NATW)*; IEEE: New York, NY, USA, 2013.
13. Weiss, E.L.; Efrat, Z.; Givati, G. System and Method for Detection of Counterfeit and Cyber Electronic Components. U.S. 12,105,857 B2, 1 October 2024.
14. Poncheri, A. AS6081-Fraudulent/Counterfeit Electronic Parts; Avoidance, Detection, Mitigation and Disposition-Distributors. In *Counterfeit Electronics Symposium*; Surface Mount Technology Association: Eden Prairie, MN, USA, 2012.
15. AS5553; Counterfeit Electronic Parts; Avoidance, Detection, Mitigation and Disposition. SAE International Standards: Warrendale, PA, USA, 2009. Available online: www.sae.org (accessed on 24 August 2026).
16. AS5553A; Fraudulent/Counterfeit Electronic Parts; Avoidance, Detection, Mitigation, and Disposition. SAE International Standards: Warrendale, PA, USA, 2013.
17. Weiss, E. Secure Hardware Assurance Using Visual AI on AOI Imaging of Electronic Assemblies. *J. Hardw. Syst. Secur.* **2026**, *10*, 3. [\[CrossRef\]](#)
18. Weiss, E. AI Detection of Body Defects and Corrosion on Leads in Electronic Components, and a study of their Occurrence. In *2022 IEEE International Symposium on the Physical and Failure Analysis of Integrated Circuits (IPFA)*; IEEE: New York, NY, USA, 2022; pp. 1–6.
19. Weiss, E. Revealing Hidden Defects in Electronic Components with an AI-Based Inspection Method: A Corrosion Case Study. *IEEE Trans. Compon. Packag. Manuf. Technol.* **2023**, *13*, 1078–1080. [\[CrossRef\]](#)
20. Weiss, E. Electronic Components Provenance Assurance Through AI and Big Data in Assembly Processes. *IEEE Trans. Compon. Packag. Manuf. Technol.* **2024**, *14*, 2107–2115. [\[CrossRef\]](#)
21. Zhao, W.; Gurudu, S.R.; Taheri, S.; Ghosh, S.; Sathiaselalan, M.A.M.; Asadizanjani, N. PCB Component Detection Using Computer Vision for Hardware Assurance. *Big Data Cogn. Comput.* **2022**, *6*, 39. [\[CrossRef\]](#)
22. Sathiaselalan, M.A.M.; Paradis, O.P.; Taheri, S.; Asadizanjani, N. Why is deep learning challenging for printed circuit board (Pcb) component recognition and how can we address it? *Cryptography* **2021**, *5*, 9. [\[CrossRef\]](#)
23. Li, Y.; Wang, X.; He, Z.; Wang, Z.; Cheng, K.; Ding, S.; Fan, Y.; Li, X.; Niu, Y.; Xiao, S.; et al. Industry-Oriented Detection Method of PCBA Defects Using Semantic Segmentation Models. *IEEE/CAA J. Autom. Sin.* **2024**, *11*, 1438–1446. [\[CrossRef\]](#)
24. Liu, J.; Xie, G.; Wang, J.; Li, S.; Wang, C.; Zheng, F.; Jin, Y. Deep Industrial Image Anomaly Detection: A Survey. *Mach. Intell. Res.* **2024**, *21*, 104–135. [\[CrossRef\]](#)
25. Yao, N.; Zhao, Y.; Kong, S.G.; Guo, Y. PCB defect detection with self-supervised learning of local image patches. *Measurement* **2023**, *222*, 113611. [\[CrossRef\]](#)
26. Tong, X.; Ziao, Y.; Tian, X.; Ge, H.; Wang, X. Improving accuracy of automatic optical inspection with machine learning. *Front. Comput. Sci.* **2022**, *16*, 161310. [\[CrossRef\]](#)

27. Lee, S.J. Few-Shot Adaptation of Foundation Vision Models for PCB Defect Inspection. *J. Imaging* **2025**, *11*, 415. [[CrossRef](#)]
28. Yin, Z.; Qin, L.; Zhang, F.; Yu, Y.; Shi, X.; Han, G.; Bi, Y. A feature-aware attention selection network for anomaly detection on printed circuit boards. *Eng. Appl. Artif. Intell.* **2026**, *171*, 114303. [[CrossRef](#)]
29. Pandey, S.; Singh, R.K.; Gunasekaran, A.; Kaushik, A. Cyber security risks in globalized supply chains: Conceptual framework. *J. Glob. Oper. Strateg. Ssourc.* **2020**, *13*, 103–128. [[CrossRef](#)]
30. Ghadge, A.; Weiß, M.; Caldwell, N.D.; Wilding, R. Managing cyber risk in supply chains: A review and research agenda. *Supply Chain Manag. Int. J.* **2020**, *25*, 223–240. [[CrossRef](#)]
31. Boyens, J.M. *Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. [[CrossRef](#)]
32. Mehta, D.; Lu, H.; Paradis, O.P.; Mukhil Azhagan, M.S.; Tanjidur Rahman, M.; Iskander, Y.; Chawla, P.; Woodard, D.L.; Tehranipoor, M.; Asadizanjani, N. The big hack explained: Detection and prevention of PCB supply chain implants. *ACM J. Emerg. Technol. Comput. Syst. (JETC)* **2020**, *16*, 1–25.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.