BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

AI-BASED DETECTION AND CLASSIFICATION OF PROTECTIVE COATING DEFECTS IN DEFENCE PLATFORMS : *A Standards-Aligned Survey and Proposed Framework*

Saee Sachin Kadam

Student, Department of Artificial Intelligence and Data Science
Ajeenkya DY Patil School of Engineering, Pune , India

Prof.Ashima Mahendra

Guide,Department of Artificial Intelligence and Data Science
Ajeenkya DY Patil School of Engineering, Pune, India

Abstract

Protective coatings are the first line of defence against corrosion on naval, aerospace and land defence platforms, yet their condition is still judged largely by human visual inspection, which is subjective, difficult to audit and hard to scale across a fleet. Recent deep-learning studies automate corrosion and coating-defect detection, but the systems reviewed here mostly return generic labels such as “rust” rather than the graded severity ratings that quality systems expect under ISO 4628, SSPC-VIS and AMPP/NACE practice. This paper reviews ten works published between 2018 and 2025, covering supervised CNNs, post-hoc explainability, few-shot learning, edge prognostics and acoustic-emission classification, and compares them in a year-wise survey matrix. The matrix exposes four gaps: no standards-referenced severity grading, data scarcity in defence archives, little multi-modal fusion, and security constraints that rule out centralised training. In response, we formalise a defect taxonomy mapped to the three standards families, propose a severity-scoring formulation based on extent, density and trend, and outline a five-module conceptual framework that combines multi-modal fusion, few-shot learning, standards mapping, clause-linked explanation and federated edge deployment. The framework is not implemented or evaluated here; a protocol for its evaluation is proposed as future work.

Keywords: Coating defect detection, Defence platforms, ISO 4628, SSPC-VIS, Explainable AI, Few-shot learning, Federated learning, Multi-modal fusion, Literature survey

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Introduction

Ships, aircraft and ground vehicles operate in salt-laden, humid, abrasive and thermally cycling environments, and the protective coatings on them are progressively degraded by blistering, rusting, cracking, flaking, chalking and pitting. Once a coating breaks down, the substrate corrodes, so the condition of the coating drives maintenance planning, and through it fleet readiness and safety.

Inspection in defence is not only a technical activity but a compliance activity. Quality management systems such as AS9100 require evidence that is objective, auditable and traceable, and that evidence is expressed against published standards: the ISO 4628 series for rating coating degradation (ISO 4628), ISO 12944 for corrosivity categories and paint systems (ISO 12944), and the SSPC-VIS photographic references now maintained by AMPP (SSPC-VIS). In practice, inspectors compare a surface against these references by eye. The result depends on the inspector, differs between sites and inspection dates, and is difficult to process at fleet scale as the volume of imagery from modern assets grows.

Artificial intelligence offers a way to make this process faster and more consistent, and the literature on deep-learning-based corrosion and coating-defect detection has grown quickly (Lu, 2020; Nahar et al., 2025; Wang et al., 2024;). However, the works surveyed in this paper mostly stop at detection: they output a class label or a segmentation mask, not the ordinal severity grade that a quality system needs (for example ISO 4628-3, Ri 4). Their explanations, where present, are not tied to clauses of a standard. They typically assume large labelled datasets, whereas defence archives are small, imbalanced and sensitive, and they assume that imagery can be pooled centrally, which security rules often forbid.

This paper therefore asks a focused question: *what would an AI inspection system have to look like to report coating defects in the grade language of the standards, under the data and security constraints of defence?* Our contributions are:

- A structured review of ten works (2018–2025) and a year-wise literature survey matrix comparing methodology, key findings and limitations (Literature Review section).
- A defect taxonomy for protective coatings mapped to ISO 4628, SSPC-VIS and AMPP/NACE (Standards Landscape and Defect Taxonomy section).
- A gap analysis that identifies four unmet requirements (Research Gaps and Challenges section).
- A severity-scoring formulation and a five-module conceptual framework that address those gaps, together with a proposed evaluation protocol (Proposed Conceptual Framework and Proposed Evaluation Protocol sections).

The work is deliberately limited to literature survey and gap analysis. No model is trained, no classified data is used and no field trial is conducted; the Limitations and Scope section states the consequences of this scope.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Standards Landscape and Defect Taxonomy

Three Families of Inspection Standards

ISO. The ISO 4628 series defines how the degradation of a coating is designated and rated, with separate parts for blistering (Part 2), rusting (Part 3), cracking (Part 4), flaking (Part 5), chalking (Part 6), delamination and corrosion around a scribe (Part 8) and filiform corrosion (Part 10), under a general introduction in Part 1 (ISO 4628). Most parts rate the quantity and, where relevant, the size of a defect on a 0–5 scale. ISO 12944 addresses a different question, classifying the corrosivity of the service environment and guiding the choice of protective paint system for steel structures (ISO 12944).

SSPC. The SSPC-VIS series supplies photographic references for surface preparation and for the degree of rusting on painted steel (SSPC-VIS). Because the judgement is made against pictures, it is the most direct target for computer vision.

AMPP (formerly NACE/SSPC). AMPP was formed from the merger of NACE International and SSPC and continues their corrosion-control and coating-inspection practices. Inspection procedures published under this name are the third reference an inspection system may need to satisfy.

The three families use different nomenclature, rating scales and photographic references. A practical AI tool must therefore be able to express a defect in whichever grade system the receiving quality system requires, and ideally cross-map between them.

Defect Taxonomy

Table 1 lists the six defect classes considered in this work and the ISO 4628 part and rated attribute for each. Pitting has no dedicated ISO 4628 part and would have to be mapped to AMPP/NACE-style corrosion assessment practice; this mapping is proposed here and should be confirmed against the current editions of the standards.

Table 1: Defect taxonomy mapped to ISO 4628

Defect	ISO 4628	Attribute rated	Scale
Blistering	Part 2	Quantity and size of blisters	0–5 each
Rusting	Part 3	Percentage of surface rusted	Ri 0–Ri 5
Cracking	Part 4	Quantity and size of cracks	0–5 each
Flaking	Part 5	Quantity and size of flaked area	0–5
Chalking	Part 6	Powder transferred to tape or cloth	0–5

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Defect	ISO 4628	Attribute rated	Scale
Pitting	None	Pit density and depth (AMPP/NACE mapping, proposed)	To be defined

Rating attributes are summarised from the standards for orientation; the standards remain the authoritative source.

Literature Review

Method

Papers were collected from IEEE Xplore, ScienceDirect and Springer, supplemented by open-access and preprint venues, for the topics of AI-based coating and corrosion defect detection, explainable AI for defect models, self-healing coating prognostics and few-shot inspection. Ten works published between 2018 and 2025 were retained: seven primary studies that propose or evaluate a method (Liu et al., 2018; Lu, 2020; Hoseini et al., 2022; Zhang et al., 2023; Rahman et al., 2023; Wang et al., 2024; Pasupuleti, 2025), and three reviews (Nahar et al., 2025; For each primary study we extracted the methodology, the key finding and the stated or evident limitation, and compared them in Table 2. Studies are ordered by year.

Survey Matrix

Table 2: Year-wise literature survey matrix

Author, year	Methodology and key findings	Limitations
Liu *et al.*, 2018	Faster R-CNN and VGG19 with TLCAF feature extraction and an SVM classifier, integrated with drone and infrared thermography; 89.54% accuracy in classifying coating breakdown and corrosion types.	No standards-referenced severity grading. Not designed for low-label or federated defence scenarios.
Lu, 2020	Multi-layer CNN classifying image patches, trained with the Adam optimiser; shows CNNs are feasible for coating-defect classification.	No severity grading, no explainability, not suited to small or imbalanced datasets.
Hoseini *et al.*, 2022	Few-shot prototypical networks on 3D height profiles from laser-scanning microscopy; results similar to conventional CNNs with far fewer annotated samples.	Outputs a general-purpose quality metric, not a standard-compliant grade. No explainability or federated learning.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Author, year	Methodology and key findings	Limitations
Zhang *et al.*, 2023	Post-hoc explainability: morphological defect characteristics fed to decision-tree ensembles, giving feature-level explanations such as boundary irregularity.	Explanations are not tied to inspection standards such as ISO 4628. Not for low-label or federated settings.
Rahman *et al.*, 2023	Decision forest on features from acoustic-emission signals; 99.48% accuracy predicting failure severity of coating disbondment.	No image-based inspection. Uses a custom severity scale, not AMPP, NACE, SSPC or ISO.
Wang *et al.*, 2024	CNN-based corrosion recognition on steel structures; data quantity and quality affect accuracy, and instance segmentation is more sensitive to small defects.	No standards-referenced grading scheme and no explainability layer.
Pasupuleti, 2025	Edge Tensor-AI with model-predictive control; tensor features from electrochemical impedance data give edge-deployable prognostics for self-healing coatings.	Addresses continuous prognosis, not image-based defect detection or standards-referenced grading.

Lu (2020) is a patent application; the remaining entries are journal, conference or preprint papers.

Synthesis

Supervised detection. The image-based studies (Liu et al., 2018; Lu, 2020; Wang et al., 2024) establish that CNN and R-CNN families can detect and classify coating breakdown and corrosion, and Wang et al. (2024) show that accuracy is tied to the quantity and quality of training data and that segmentation copes better with small defects than box-level detection. All three end at a defect label, not at a rated grade.

Data-efficient learning. Hoseini et al. (2022) are the only study in the set that tackles label scarcity directly, using prototypical networks to match CNN performance with far fewer annotated samples. Their input is a laser-scanned height profile rather than a photograph, and their output is a generic quality metric.

Explainability. Zhang et al. (2023) explain defect classifiers through morphological characteristics and tree ensembles. The explanations are useful to a data scientist but are not phrased in the vocabulary of a standard, so an auditor cannot trace a prediction to a clause.

Non-image sensing. Pasupuleti (2025) and Rahman et al. (2023) use electrochemical impedance and acoustic-emission signals respectively. They show that coating condition is observable through channels

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

other than vision, which motivates fusion, but each produces its own prognosis or severity scale rather than a standard grade.

Reviews. Nahar et al. (2025) taxonomise supervised, unsupervised and emerging deep-learning approaches to industrial defect detection and confirm that supervised models dominate yet need large labelled datasets; standards compliance and defence-specific constraints are outside their scope. The reviews of deep learning in corrosion assessment and of image recognition in corrosion protection provide broader context on techniques and challenges, and are used here as background rather than as primary comparators.

Table 3 condenses the matrix into a capability comparison. Its entries are derived from the limitations recorded in Table 2 and from the study descriptions, and "Yes" means the capability is addressed, "Part" that it is addressed in part, and "No" that it is not addressed.

Table 3: Capability comparison of primary studies

Study	Vis.	Std.	FS	MM	XAI	Edge
Liu et al. (2018)	Yes	No	No	Part	No	No
Lu (2020)	Yes	No	No	No	No	No
Hoseini et al. (2022)	Yes	No	Yes	No	No	No
Zhang et al. (2023)	Yes	No	No	No	Yes	No
Rahman et al. (2023)	No	No	No	No	No	No
Wang et al. (2024)	Yes	No	No	No	No	No
Pasupuleti (2025)	No	No	No	No	No	Yes
Proposed framework	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**

Vis.: vision-based input (image or surface profile); Std.: standards-graded severity; FS: few-shot or data-efficient learning; MM: multi-modal fusion; XAI: explainability; Edge: edge or federated deployment. The last row is a design target, not a demonstrated result.

Research Gaps and Challenges

Reading Tables 2 and 3 together, five of the seven primary studies work from vision-based input, but none produces a standards-graded severity output, and each of few-shot learning, explainability and edge deployment appears in exactly one study (Figure 1). No single study combines the capabilities. Four gaps follow.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

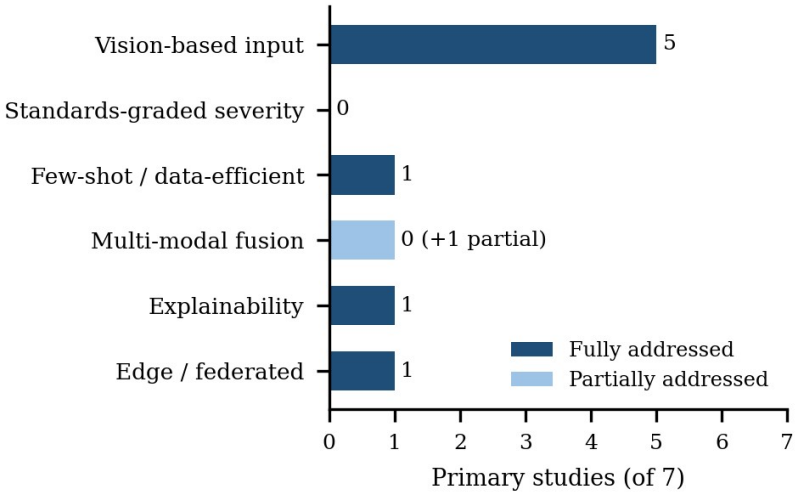


Figure 1: Number of the seven primary studies that address each capability, derived from Table 3.

Gap 1: standards-compliant grading. Existing models emit generic labels such as "rust" rather than a severity grade such as ISO 4628-3 Ri 4. The challenge is to train models that produce ordinal grades aligned with several standard families at once, and to respect the fact that being one grade wrong is less serious than being four grades wrong.

Gap 2: data scarcity in defence. Defence archives are small, imbalanced and sensitive, which makes large-scale supervised learning impractical (Nahar et al., 2025; Wang et al., 2024). Data-efficient techniques such as few-shot and meta-learning are needed, but the only such study reviewed (Hoseini et al., 2022) does not output standard grades.

Gap 3: multi-modal fusion for robustness. Most systems rely on single-modality RGB data, which can be misled by shadows, dirt and surface texture. Sensing channels that carry independent evidence, such as 3D height profiles (Hoseini et al., 2022), infrared thermography (Liu et al., 2018), electrochemical impedance (Pasupuleti, 2025) and acoustic emission (Rahman et al., 2023), are studied separately. The challenge is to fuse visual, 3D and electrochemical data into one assessment.

Gap 4: security and data governance. Centralising raw imagery from defence sites is often not permitted. Models must run at the edge and be trainable across sites without sharing raw data. Only Pasupuleti (2025) considers edge deployment, and none of the reviewed works considers federated learning.

A fifth, cross-cutting requirement is *explainability tied to clauses*: an inspector or auditor must be able to see which part of the surface and which rule of which standard produced a grade. The feature-level explanations of Zhang et al. (2023) are a starting point but do not reach this level.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Proposed Conceptual Framework

Severity-Scoring Formulation

To convert raw inspection observations into a discretised grade, we propose scoring each defect class c observed on an inspected region using three measurements: extent, density and trend.

$$E = A_d / A_{total}, \quad D = N_d / A_{total}, \quad T = (E_t - E_{t-1}) / \Delta t \tag{1}$$

Here the extent E is the defected area as a fraction of the total inspected area, the density D is the number of discrete defects (blisters, cracks, pits) per unit inspected area, and the trend T is the change in extent between two inspections divided by the interval Δt. Trend uses the previous inspection of the same region; when none exists, T is set to zero and the remaining weights are renormalised. The three terms are normalised and combined into a score

$$S = w_E \varphi_E(E) + w_D \varphi_D(D) + w_T \varphi_T(T) \tag{2}$$

where the three weights are non-negative and sum to one, and each φ maps its input to [0, 1], so S lies in [0, 1]. The score is then discretised into a grade band by ordered thresholds,

$$g = k \ \ if \ \ \tau_k \leq S < \tau_{k+1}, \quad k = 0, 1, \ldots, 5 \tag{3}$$

and a standards mapper M translates the pair (class, g) into an ISO 4628 rating, an SSPC-VIS grade and an AMPP/NACE category. Where a standard defines its rating purely by area, as ISO 4628-3 does for rusting (Table 4), the mapper applies the standard's own thresholds to E directly, and S serves as the cross-standard and prioritisation score. The weights w and thresholds τ are parameters to be calibrated with expert-labelled data; this paper does not set them.

Table 4: ISO 4628-3 rusting grades (nominal area)

Grade	Nominal rusted area
Ri 0	0%
Ri 1	0.05%
Ri 2	0.5%
Ri 3	1%
Ri 4	8%
Ri 5	about 40–50%

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Values as commonly tabulated for ISO 4628-3, given for orientation and to be verified against the standard before use.

Architecture

The framework has five modules (Figure 2), each aimed at one or more of the gaps identified in the Research Gaps and Challenges section.

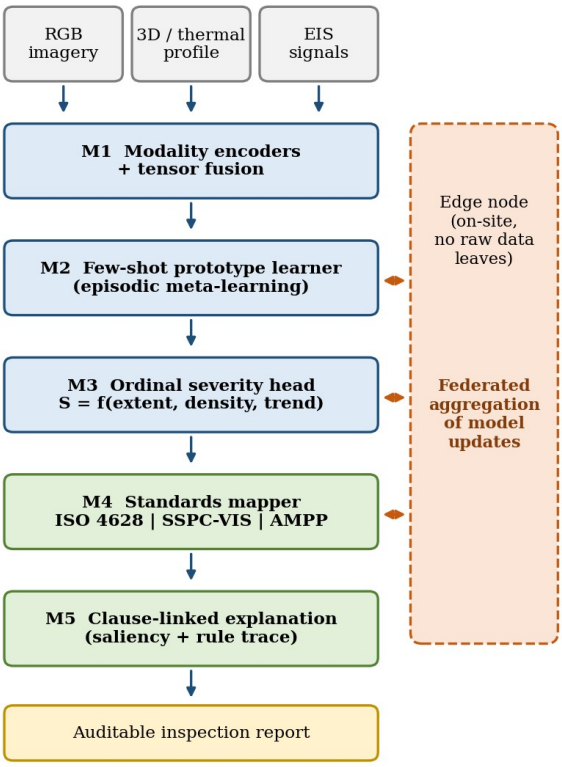


Figure 2: Conceptual framework. Modules M1–M3 (blue) produce a severity estimate, M4–M5 (green) make it standards-compliant and explainable, and the federated edge loop (orange) keeps raw data on site.

M1: modality encoders and tensor fusion (Gap 3). Separate encoders process RGB imagery, 3D or thermal profiles and, where available, electrochemical impedance features; their outputs are fused into a joint tensor representation. Fusion is intended to make the assessment robust to conditions that fool a single modality, and the module must tolerate a missing modality, since not every site will hold every data type.

M2: few-shot prototype learner (Gap 2). A prototype-based meta-learner is trained over episodes sampled by defect class and grade, so a new grade or defect type can be learned from a handful of labelled examples. This is the most direct answer to small and imbalanced defence archives, and follows the direction shown by Hoseini et al. (2022).

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

M3: ordinal severity head (Gap 1). Instead of independent class outputs, an ordinal head with a rank-consistent loss predicts the grade so that near misses cost less than distant ones. It computes E, D and T from the detected regions and produces S via Eq. (2).

M4: standards mapper (Gap 1). A rule table implements M, translating (class, g) into ISO 4628, SSPC-VIS and AMPP/NACE grades and cross-mapping between them, so the same finding can be reported in whichever system the receiving quality system uses. Because the table is explicit, it can be reviewed and updated by inspectors when a standard is revised, without retraining the model.

M5: clause-linked explanation (cross-cutting requirement). The explanation has two layers: a visual layer that highlights the regions driving the prediction, for example with gradient-based saliency, and a rule-trace layer that states the clause and threshold applied. An illustrative trace reads: *"Rusted area is about 1% of the inspected surface, which corresponds to ISO 4628-3 Ri 3."* Each report is stored with its inputs, module versions and trace so it can be audited.

Federated edge deployment (Gap 4). Inference runs on on-site hardware. Each site fine-tunes on its own archive and shares only model updates for aggregation, in the manner of federated averaging, so raw imagery never leaves the site. Open issues are that data differ between sites and that model updates can leak information, for which secure aggregation and differential privacy are candidate mitigations.

Proposed Evaluation Protocol

The framework has not been built, so no results are reported. To keep the claims testable, we propose the following protocol for the implementation stage.

- **Grading quality:** ordinal mean absolute error and quadratic weighted kappa against expert grades, plus macro-F1 for defect class and a per-grade confusion matrix.
- **Data efficiency:** N-way K-shot accuracy as K is reduced, compared with a fully supervised CNN trained on the same images.
- **Fusion benefit:** performance with each modality removed in turn, including on images with shadows or surface staining.
- **Explanation quality:** agreement between the clause cited by the system and the rationale given by a certified inspector, and fidelity tests that remove highlighted regions and observe the change in grade.
- **Federated behaviour:** accuracy against a centralised baseline, communication rounds required, and robustness when sites hold different defect distributions.
- **Human agreement:** inter-rater agreement between the system and inspectors, compared with agreement among the inspectors themselves.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Because classified data will not be accessible, initial experiments would use public corrosion and coating image sets and laboratory panels rated by experts, with the defence setting emulated through subsampling and site partitioning.

Limitations and Scope

The paper reviews ten works, which is enough to expose recurring gaps but is not an exhaustive systematic review; extending the corpus is a natural next step. The capability comparison in Table 3 is based on the published descriptions and may miss features that a paper implements but does not emphasise. The framework is conceptual: the weights and thresholds in Eqs. (2) and (3) are uncalibrated, the standards mapping is a proposed interpretation that should be validated by certified inspectors against current standard editions, and no experiment supports any performance claim. Most reviewed studies address marine, civil or industrial settings, so their findings may not transfer directly to defence coatings. Classified data, model training and field deployment are outside the scope of this stage.

Conclusion and Future Work

Manual visual inspection of protective coatings on defence platforms is subjective, inconsistent and difficult to audit, while the AI literature reviewed here largely stops at detection. Comparing ten works showed that none of the seven primary studies outputs standards-referenced severity grades, and that few-shot learning, explainability and edge deployment each appear in only one of them. We addressed this with a defect taxonomy mapped to ISO 4628, SSPC-VIS and AMPP/NACE, a severity formulation based on extent, density and trend, and a five-module framework that combines multi-modal fusion, few-shot learning, standards mapping, clause-linked explanation and federated edge deployment.

The next stage is implementation, beginning with the multi-modal preprocessing and tensor fusion module (M1), followed by the few-shot learner and ordinal head, calibration of the severity weights and thresholds with expert-labelled data, and evaluation under the protocol in the Proposed Evaluation Protocol section.

References

1. Association for Materials Protection and Performance. (n.d.). SSPC-VIS series: Visual standards for abrasive blast cleaned steel and for painted steel surfaces.
2. Hoseini, S., et al. (2022). Automated defect detection for coatings via height profiles obtained by laser-scanning microscopy. Machine Learning with Applications, 10.
3. International Organization for Standardization. (n.d.). ISO 12944: Paints and varnishes — Corrosion protection of steel structures by protective paint systems.
4. International Organization for Standardization. (n.d.). ISO 4628 series: Paints and varnishes — Evaluation of degradation of coatings.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

5. Liu, L., Tan, E., Cai, Z. Q., Zhen, Y., & Yin, X. J. (2018). An integrated coating inspection system for marine and offshore corrosion management. In Proceedings of ICARCV.

6. Lu. (2020). Coating defect detection method based on deep learning (Chinese Patent Application CN 112150417 A).

7. Nahar, L., et al. (2025). AI-enabled defect detection in industrial products: A comprehensive survey, key insights and future research challenges. Advanced Engineering Informatics, 69.

8. Pasupuleti, M. K. (2025). Edge Tensor-AI enabled predictive control of self-healing anti-corrosion coatings for marine robots. IJAIRI, 5(12), 291–313.

9. Rahman, N. A. A., May, Z., Jaffari, R., & Hanif, M. (2023). Failure severity prediction for protective-coating disbondment via the classification of acoustic emission signals. Sensors, 23(15).

10. SAE International. (n.d.). AS9100: Quality management systems — Requirements for aviation, space and defense organizations.

11. Wang, J., et al. (2024). Recognition of steel bridge corrosion considering the effect of training dataset quality. Case Studies in Construction Materials, 21.

12. Zhang, J., Cosma, G., Bugby, S., & Watkins, J. (2023). ForestMonkey: Toolkit for reasoning with AI-based defect detection and classification models. arXiv:2307.13815.