AI-empowered 3D X-ray analysis of solder joint fatigue after board level vibration testing

Amir Ghorbani Ghezeljehmeidan
Department of Microelectronics
Delft University of Technology
Delft, The Netherlands
amir.ghorbani@tudelft.nl

Varun Thukral
NXP Semiconductors
Nijmegen, The Netherlands

Fei Xu
NXP Semiconductors
Nijmegen, The Netherlands

Willem D. van Driel
Department of Microelectronics
Delft University of Technology
Delft, The Netherlands

2025 IEEE 27th Electronics Packaging Technology Conference (EPTC) | 979-8-3315-6145-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/EPTC67330.2025.11392644

***Abstract*—Mission-critical electronic systems demand early and accurate detection of solder joint degradation to ensure reliability. Quad Flat No-Lead (QFN) packages, widely used in automotive and industrial applications, are especially prone to vibration-induced solder fatigue. However, traditional failure analysis methods (e.g., dye-and-pry, cross-sectioning, manual X-ray inspection) are labor-intensive and often insufficient to detect early-stage cracks. This paper presents an automated inspection framework that combines high-resolution 3D X-ray tomography with a YOLOv11-based deep learning model to detect and segment vibration-induced cracks in QFN solder joints. The pipeline achieves precise localization of cracks in volumetric data, discriminates them from voids, and extracts morphological descriptors through parametric fitting. By statistically correlating these image-derived crack features with electrical resistance measurements recorded in situ during vibration tests, we establish a direct link between physical crack evolution and functional degradation of the joint. The results demonstrate that our AI-driven method can automatically identify tiny solder cracks and reliably offer predict impending interconnect failures in comparable granularity of traditional inspection techniques, surpassing them in speed. This approach offers a powerful prognostic health monitoring tool for electronic packaging, and it is extensible to other package types and stress conditions.**

***Index Terms*—Board-Level Reliability, Solder Joint Cracks, 3D X-ray Tomography, Deep Learning, YOLO, Prognostics and Health Monitoring (PHM), Vibration Testing, QFN Package.**

I. Introduction

Ensuring the reliability of electronic systems is critical, especially for safety- and mission-critical applications such as automotive electronics. A significant portion of electronics failures can be traced to solder interconnect degradation [1], with mechanical vibrations identified as a major contributor (around 20% of failures [2]). In particular, Quad Flat No-Lead (QFN) packages, popular in automotive and industrial domains, are known to be vulnerable to vibration-induced damage in their solder joints [3]–[5]. Repeated vibrational stress causes PCB flexure and solder fatigue, potentially initiating micro-cracks at the solder-pad interface. If undetected, these cracks can propagate and eventually lead to open circuits or intermittent connections.

Conventional failure analysis techniques for solder joints have notable limitations. Methods like dye-and-pry, cross-sectioning, or manual 2D/3D X-ray inspection require significant labor and often destructively prepare the sample, yet still may miss fine early-stage cracks. Moreover, manual inspections are subjective and lack the consistency and resolution needed for prognostic health monitoring. Prior studies have attempted to monitor solder joint health via electrical signals, for instance, by tracking changes in joint resistance as a proxy for crack growth [6], [7]. It is demonstrated that serial ohmic resistance measurements can serve as a structural health monitoring tool for QFN solder joints under fatigue. While such electrical monitoring techniques detect when a failure is imminent or has occurred, they do not directly localize or quantify the physical damage.

In recent years, the electronics reliability community has started to explore data-driven and AI-assisted approaches to overcome these challenges. Machine learning algorithms have been applied to predict solder joint failures and interpret complex degradation data, showing improved prognostic accuracy over traditional methods. For instance,recent surveys show the application of AI techniques for failure analysis in industrial systems, underscoring their ability to handle heterogeneous data and detect subtle patterns [8]. Other works have combined finite element simulations with neural networks to estimate solder fatigue life, or employed deep learning on in-situ sensor data for condition monitoring of solder joints [9]. More directly relevant to this work, researchers have begun applying computer vision and deep learning to X-ray images of solder joints [4]. These advances motivate the application of modern computer vision techniques to enhance solder joint inspection and prognostics.

This paper introduces a novel AI-driven inspection pipeline that addresses the need for early and precise crack detection in solder joints. Our approach specifically targets QFN package solder joints subjected to board-level vibration tests. We integrate high-resolution 3D X-ray imaging with a state-of-the-art YOLOv11 deep learning model to automatically detect and segment crack regions. The system then performs morphological analysis on the identified cracks (e.g., measuring crack length, area, and curvature) and correlates these metrics with real-time electrical resistance data recorded during vibration testing. By linking physical damage progression with electrical performance degradation, the proposed method provides a holistic perspective of solder joint health. In the following sections, we describe the experimental setup and AI

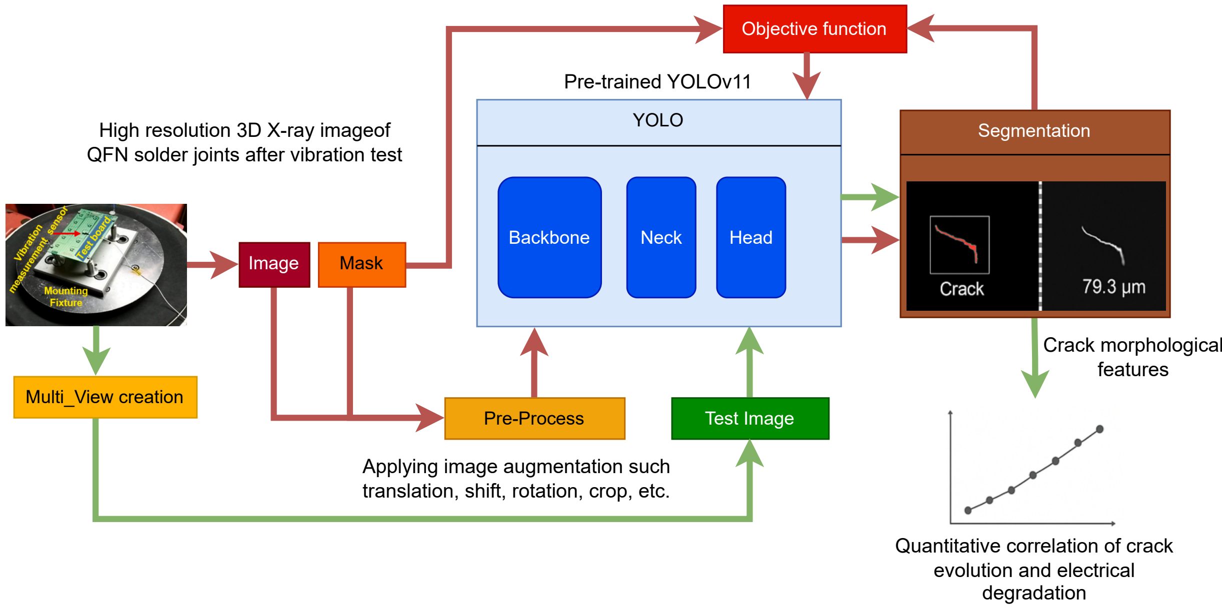

Fig. 1. Workflow for automated crack detection and analysis using 3D X-ray images of QFN solder joints post-vibration testing. A pre-trained YOLOv11 model performs defect localization, aided by image-mask pairs and augmentation. Segmented cracks are fitted to extract morphological features, which are quantitatively correlated with electrical degradation.

methodology, present the detection and analysis results, and discuss the implications for prognostic health monitoring of electronic packages.

II. Methodology

A. *Vibration Test and Data Acquisition*

Board-level vibration testing was conducted to induce solder joint fatigue in QFN assemblies under controlled conditions [10]. The test vehicle consisted of a standard PCB onto which QFN components were soldered. Each QFN package measured 9×9 mm and had 56 pins, with an internal daisy-chain design to facilitate resistance monitoring at specific joints. In particular, the QFNs were instrumented with 4-wire sense circuitry on selected pins, allowing precise measurement of solder interconnect resistance in situ. The test board followed a JEDEC-compliant layout for vibration/drop tests and accommodated three QFN components, enabling simultaneous monitoring of 24 solder joints (eight corner joints across the three packages).

The vibration profile was a swept-sine excitation from 20 Hz to 2 kHz (spanning the board's fundamental modes) at a controlled acceleration amplitude of approximately $3g_{\mathrm{RMS}}$. The board was mounted on an electrodynamic shaker via a rigid fixture (Figure 1), and an accelerometer on the fixture provided feedback to maintain the target acceleration. All tests were performed at ambient room temperature (20°C) to eliminate thermal influences. The vibration was applied until solder joint failure occurred, as indicated by a change in the electrical resistance of the monitored joints.

During testing, a data acquisition system recorded the resistance of each instrumented joint in real time using a 4-wire measurement setup and high-precision source measurement units (SMUs). The measurement system sampled each channel at 100 Hz. To capture failure events, a threshold-based logging strategy was employed: only significant resistance "drops" or spikes (indicative of an open circuit or sudden increase in joint resistance beyond a predefined threshold) were recorded as failure events for each joint. In other words, the system focused on logging the moments of abrupt resistance change ("drop values") rather than storing continuous time-series data, which reduced data volume and highlighted the exact moment of electrical failure. This approach is consistent with prior board-level vibration studies that monitor for the instant of interconnect rupture [3]. For each failed solder joint, the final resistance jump (to an open-circuit or high-resistance state) was thus captured, providing a clear endpoint for that joint's lifetime. The number of vibration cycles endured by each joint before failure was also noted via the shaker's controller logs.

After completing the vibration tests, the failed QFN samples were non-destructively examined using 3D X-ray microscopy. A high-resolution X-ray tomography system was used to scan the region around each QFN's solder joints. The X-ray system generated a series of 2D projection images (radiographs) at incremental rotation angles, which were reconstructed into a volumetric image (3D tomography) of the soldered connections. In our setup, we obtained over 1000 cross-sectional slice images per device. These slices have a voxel size on the order of a few micrometers, sufficient to resolve small crack gaps within the solder material. The reconstructed 3D volume clearly shows the solder joints, PCB pad interfaces, and any internal defects such as cracks or voids (see Fig. 3

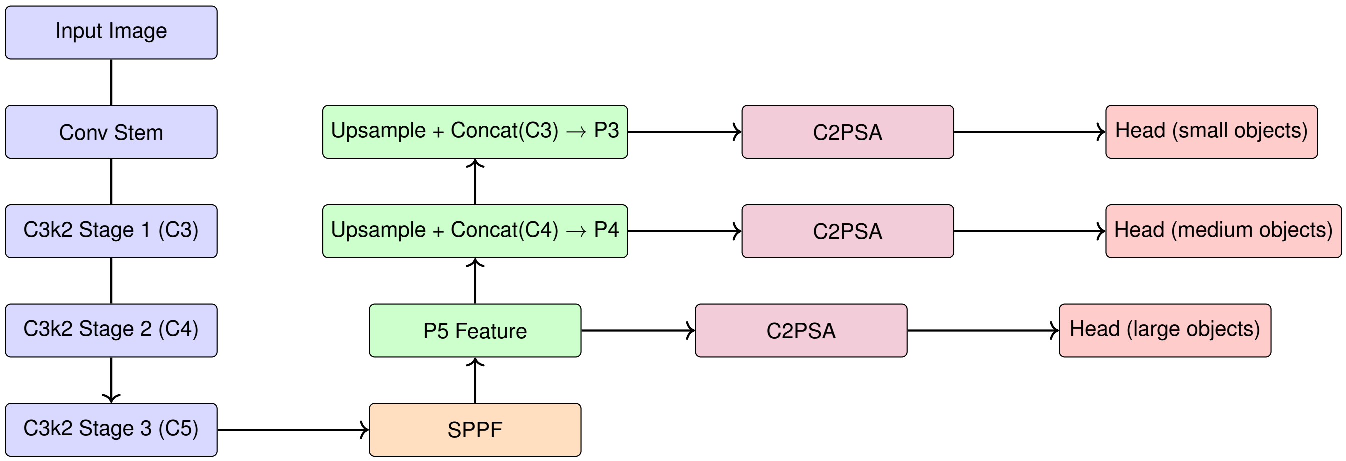

Fig. 2. YOLOv11 architecture overview. The backbone (blue) extracts features via C3k2 blocks. An SPPF module (orange) compresses context. The neck (green) fuses features into pyramid levels (P3–P5). C2PSA modules (purple) enhance spatial focus before the detection heads (red) output predictions at multiple scales.

for a schematic). By inspecting the 3D X-ray volumes, we confirmed that the electrically failed joints exhibited physical cracking at or near the PCB-pad interface, as expected for vibration-induced fatigue.

B. Deep Learning-Based Crack Detection

We employ a YOLOv11-based framework for pixel-wise crack detection and segmentation [11]. YOLOv11 [12] follows the backbone–neck–head paradigm and introduces several architectural enhancements that improve both efficiency and sensitivity to small defects. The backbone is built on lightweight C3k2 blocks, which replace the C2f modules of YOLOv8, reducing computation while maintaining strong feature representation. A Spatial Pyramid Pooling–Fast (SPPF) layer at the deepest stage compresses multi-scale context at minimal cost. The neck fuses features across scales (P3–P5) through upsampling and concatenation, enabling robust multi-resolution representation of cracks with varying size and orientation. In addition, optional C2PSA (Cross-Stage Partial with Spatial Attention) modules enhance spatial focus, which is especially beneficial for capturing subtle, low-contrast solder cracks in noisy X-ray slices. Finally, the detection head produces predictions at three scales, improving localization of both fine and extended crack structures. The model is first pretrained on large public crack datasets (e.g., concrete and pavement crack imagery) and general-purpose datasets such as COCO, then fine-tuned on annotated QFN solder-joint X-ray slices. This two-stage training leverages transfer learning to mitigate the limited availability of domain-specific data, a strategy also common in structural crack detection tasks [5].

C. Data Pre-processing and Augmentation

To account for imaging variations and address limited annotated training data, we apply data augmentation on the X-ray slices prior to training. Augmentations include random rotations, small translations, horizontal/vertical flips, and intensity adjustments (brightness/contrast). Such strategies have been shown to expand the diversity of the data set and improve the robustness of the model under variable imaging conditions [13]. These augmentations ensure the network generalizes to practical inspection scenarios.

D. Multi-View Segmentation and Fusion

Since cracks are not always clearly visible from a single slice orientation, we adopt a multi-view strategy. In addition to the native XY (axial) slices of the reconstructed X-ray volume, cross-sections in XZ and YZ planes are also extracted. Each slice is processed independently by the YOLOv11 model. Multi-view approaches leverage strong 2D segmentation networks for 3D [14]. After segmentation, masks from all three orientations are projected into the 3D volume and fused into a composite crack mask. Fusion is performed by weighted averaging the binary masks across orientations, effectively combining complementary crack evidence. This approach ensures that small or orientation-dependent cracks, often missed in a single view, are retained in the final mask.

E. Morphological Feature Extraction

From the fused 3D crack masks, morphological descriptors are extracted. The most relevant include:

- **Crack length along solder-pad interface:** the projected span of the crack parallel to the interface, strongly correlated with joint degradation.
- **Crack width:** the maximum local opening of the crack within the solder volume.

These features provide objective metrics for comparing damage severity across samples, and are consistent with prior findings that interface crack length is a dominant indicator of solder joint failure under vibration stress [3], [7].

III. Results

A. Segmentation Performance

The YOLOv11-based model achieved robust performance in segmenting cracks from QFN solder joint X-ray slices. A total of 12 solder joint corners were used for model training. Each corner contributed approximately 15 manually annotated

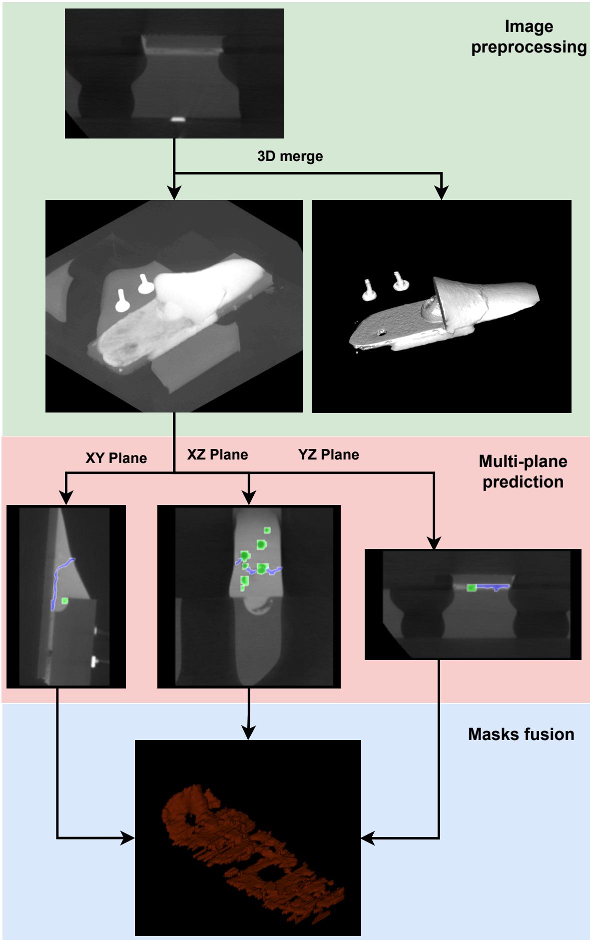

Fig. 3. Workflow of the prediction phase. Raw X-ray projections are reconstructed into a 3D volume, from which multi-view slices are extracted. For each slice, crack and void masks are generated and subsequently fused to obtain the final crack mask.

crack/void slices, resulting in a base training set of 180 annotated images. To address data scarcity and mitigate the risk of overfitting, we applied six augmentation strategies (random rotations, translations, flips, intensity adjustments, cropping, and minor geometric distortions) to each slice, yielding an expanded training dataset of over 1000 images. For validation, a separate set of 25 manually annotated slices was used without augmentation to provide an unbiased performance assessment. Manual annotation of solder cracks is inherently prone to human variability, especially when delineating very thin or low-contrast features. To mitigate this, annotations were carefully verified by multiple passes, ensuring consistent labeling across the dataset. Table I summarizes the segmentation performance of the YOLOv11 model on the annotated validation set.

TABLE I
THE FINE-TUNED YOLOv11 PERFORMANCE METRICS OVER VALIDATION DATASET.

Class	Image	Instance	Precision (%)	Recall (%)	F1 (%)	mAP@50 (%)
crack	25	47	0.81	0.71	0.75	0.76
void	17	25	0.91	0.72	0.80	0.82

$$\text{Precision} = \frac{TP}{TP + FP} \tag{1}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{2}$$

$$\text{F1-Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{3}$$

$$\text{mPA@50} = \frac{1}{N} \sum_{i=1}^{N} \frac{TP_i}{TP_i + FP_i + FN_i} \tag{4}$$

Precision measures how many of the predicted positive samples are actually correct. **Recall** measures how many of the actual positive samples are correctly detected. **F1-Score** is the harmonic mean of Precision and Recall, providing a balanced measure of both. **Mean Pixel Accuracy at 50 (mPA@50)** measures the average pixel-wise accuracy across all classes or regions, evaluated at a confidence threshold of 0.5. It reflects the overall segmentation accuracy while balancing false positives and false negatives. The model achieves a precision of 0.81 and recall of 0.71 for cracks, corresponding to a mean average precision (mAP@50) of 0.76. Although slightly lower than void segmentation (mAP@50 = 0.821), the crack detection metrics remain within the range reported for state-of-the-art crack segmentation in domains [5]. The relatively lower recall for cracks reflects the difficulty of detecting very fine or low-contrast features in 3D X-ray slices, while the high precision demonstrates that detected crack regions are reliable and rarely confused with background. These results confirm that the model provides accurate and consistent segmentation of solder joint defects, with performance sufficient for subsequent morphological analysis and correlation with electrical degradation.

B. Segmentation Across Slice Orientations

Crack visibility varied significantly across planes. To further improve the model performance in detecting the cracks. We utilize the different slices along 3 main orientations.

- In the **XY plane**, cracks aligned with the pad interface were generally visible, though some appeared as faint hairlines.
- In the **XZ and YZ planes**, vertical or oblique cracks that were nearly invisible in XY slices became clearly delineated.
- The quality of segmentation also depended on slice density: insufficient slice spacing could miss thin cracks, while denser sampling improved continuity.

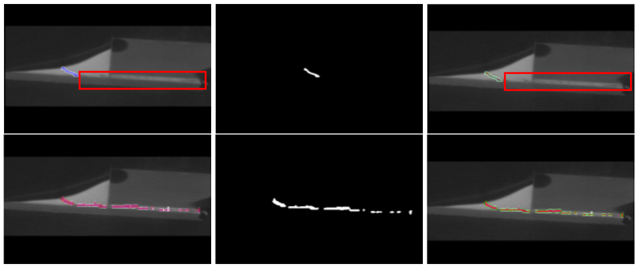

Fig. 4. Showcase of improved detection of subtle crack regions initially missed by the model. The top row shows the direct model output, while the bottom row presents the fused mask result.

Figure 4 illustrates a side-by-side comparison of YZ-plane segmentation before and after fusion, showing improved crack continuity and reduced fragmentation. This behavior aligns with prior findings that multi-view segmentation improves completeness in 3D defect detection [15].

C. Mask Fusion and Crack Volume Extraction

The three orthogonal YOLO masks (XY, XZ, YZ) are re-oriented to a common (Z, Y, X) grid via axis permutation and optional mirroring, then center-aligned by integer translations so all volumes share one canvas. Let $M_i \in \{0,1\}^{Z\times Y\times X}$ denote the aligned binary mask from view $i \in \{\text{xy,xz,yz}\}$.

To tolerate small mis-registrations and soften jagged boundaries, each M_i is converted into a distance-aware probability field $P_i \in [0,1]^{Z\times Y\times X}$ inside a tight region of interest (ROI) [16], [17]. We compute anisotropic Euclidean distance transforms to foreground/background within the ROI, using physical voxel spacings $(\Delta z, \Delta y, \Delta x)$, and form the signed distance

$$s(\mathbf{p}) \;=\; d(\mathbf{p}, \text{inside}) \;-\; d(\mathbf{p}, \text{outside}),$$

with $s > 0$ inside the mask, $s < 0$ outside. A radial kernel maps the distance to probability:

$$P_i(\mathbf{p}) \;=\; \begin{cases} 1, & s(\mathbf{p}) \geq 0, \\ \exp\big(-\, s(\mathbf{p})^2/(2\,\sigma_{\text{phys}}^2)\big), & s(\mathbf{p}) < 0, \end{cases}$$

where $\sigma_{\text{phys}} > 0$ (in *physical* units) controls tolerance to misalignment; inside-mask voxels remain at probability 1 and outside decays smoothly with distance. This distance based softening is standard in robust segmentation [18].

Per-view probabilities are combined by a weighted noisy-OR rule that preserves evidence from any view while down-weighting weak contributors. With nonnegative weights α_i (equal or validation-derived), the fused probability is

$$P_{\text{fused}}(x, y, z) \;=\; 1 - \prod_{i\in\{\text{xy,xz,yz}\}} \Big(1 - \alpha_i\, P_i(x, y, z)\Big),$$

which is order-invariant, monotone in each P_i, and reduces to a set union for binary inputs; it matches the classical noisy-OR in probabilistic models [19].

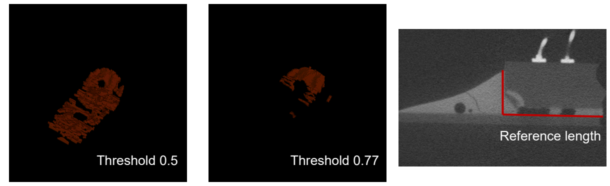

Fig. 5. Effect of applying different thresholds during mask fusion. Higher probability thresholds preserve only the most dominant crack regions. The image on the right shows the reference length used to calculate the crack length ratio.

A crack volume is obtained by thresholding P_{fused} at a user-set probability τ (Fig. 5 uses $\tau = 0.77$):

$$\hat{M} \;=\; \big\{(x, y, z) \;:\; P_{\text{fused}}(x, y, z) \geq \tau\big\}.$$

We remove small connected components (3D connectivity) and apply a light 3D binary closing with a spherical structuring element to suppress speckle and fill pinholes.

D. Discussion on Morphological Indicators

From the fused hard mask $\hat{M}$, the centerline of the crack is obtained by iterative morphological thinning; The 3D arc length is then computed by adding unit and $\sqrt{2}$ steps over connected skeleton neighbors and converted to physical units with the voxel spacing. While crack length along the solder–pad interface is a primary descriptor, length alone can overemphasize hairline, low-visibility filaments. In our pipeline, such tenuous segments are implicitly suppressed by:

1) the probability threshold τ on P_{fused}
2) volumetric post-processing (small-object removal and closing), so that only spatially consistent, high-confidence crack regions contribute to the skeleton and downstream metrics

Empirically, this yields more stable morphology (effective length) and a clearer correlation with resistance-change measurements than using raw, single-view masks. As shown in Figure 4, cracks consistently initiate at the wetting side of the joint and propagate laterally along the pad plane. Segmentation analysis quantifies the reduction in the cross-sectional area of the conductive area as the cracks grow, while concurrent in situ electrical monitoring records a corresponding increase in resistance (Figure 6). The observed increase in resistance closely correlates with the reduction in conductive cross-sectional area determined from the segmentation, such that the trend in electrical degradation quantitatively matches the evolving crack geometry [6], [7]. Importantly, at the crack ratios reported in the figure, the resistance–geometry relationship follows the same trend as previously reported in experimental and simulation studies, reinforcing the consistency of the present findings with established degradation behavior.

IV. Discussion and Conclusion

This paper presented a comprehensive AI-enabled methodology for automated solder joint crack detection and prognostics in QFN packages using 3D X-ray imaging and deep

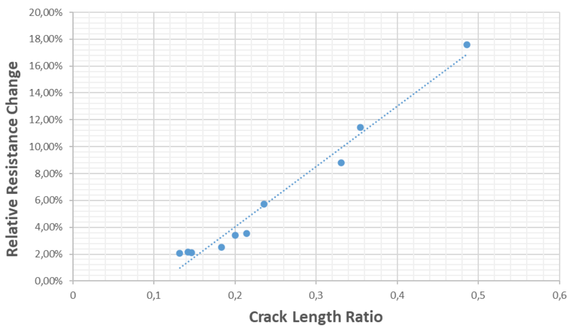

Fig. 6. Relative resistance change with respect to effective crack length ratio.

learning. By combining an advanced YOLOv11-based segmentation model with high-resolution X-ray tomography, we achieved precise automated identification of solder cracks that were previously only detectable by laborious manual analysis. The system not only finds cracks with high accuracy, but also quantifies their morphology and correlates these quantitative damage metrics with electrical degradation data. This combination of electrical and imaging information provides a direct and objective link between the evolution of physical defects and loss of functional performance.

The experimental results on vibration-tested QFN assemblies demonstrate that our approach can reliably flag early-stage cracks and predict impending failures, offering significant improvements over traditional pass/fail testing. The ability to non-destructively monitor crack growth and anticipate failures is particularly valuable for prognostic health monitoring (PHM) of electronics in safety-critical applications. Maintenance can be scheduled before catastrophic failure, and design weaknesses can be identified by analyzing where and how cracks form.

In terms of impact and future scope, the data-driven pipeline introduced here represents a shift towards more granular and automated reliability analysis in electronic packaging. While we focused on QFN packages under mechanical vibration, the methodology is inherently adaptable. The same framework could be applied to other package types such as Ball Grid Arrays (BGA) or chip-scale packages, and to other stress regimes like thermal cycling or drop shock, with minimal modifications. The deep learning model can be retrained for new defect types or imaging modalities (e.g., 2D X-ray inspection systems), making this approach broadly extensible across reliability testing scenarios.

A promising avenue for further development is the adaptation of the described pipeline into a real-time edge-capable solution. Specifically, a similar framework could be built on much lighter models to make it suitable for on-device deployment. In the broader computer-vision community there are many studies using YOLO variants [20] or Segment Anything Model (SAM) [21] for edge applications, leveraging techniques such as knowledge-distillation, student–teacher modeling, and adapter [22], modules to create domain-specialized, efficient models.

In conclusion, the integration of 3D X-ray imaging and AI-based analysis as demonstrated in this work provides a powerful tool to reveal the physics of degradation in solder interconnects. Establishing a quantitative relationship between resistance change and crack geometry provides a direct link between physical defect evolution and functional degradation of the solder joint. The experimental observations show that resistance increases systematically with crack propagation, reflecting the progressive reduction in effective conductive cross-section. By segmenting the crack morphology from high-resolution X-ray data and correlating it with in-situ resistance measurements, it becomes possible to map the electrical response to mechanical damage with high fidelity. This correlation enables non-destructive, real-time monitoring of solder joint health and offers a measurable indicator of degradation progression. The outcome of this relationship is a predictive framework in which resistance changes (in highly accurate measurement) serve as proxies for crack morphology. This allows for early detection of failure onset and quantitative estimation of remaining useful life.

Acknowledgment

The authors thank NXP Semiconductors for providing test samples and facilities.

References

[1] N. Jiang, L. Zhang, Z. Q. Liu, L. Sun, W. M. Long, P. He, M. Y. Xiong, and M. Zhao, "Reliability issues of lead-free solder joints in electronic devices," *Sci. Technol. Adv. Mater.*, vol. 20, no. 1, pp. 876–901, 2019, doi: 10.1080/14686996.2019.1640072.

[2] K. Baishya, D. M. Harvey, T. Partida Manzanera, G. Zhang, and D. R. Braden, "Failure patterns of solder joints identified through lifetime vibration tests," *Nondestructive Testing and Evaluation*, vol. 38, no. 1, pp. 147–171, 2022, doi: 10.1080/10589759.2022.2084616.

[3] F. Arabi, A. Gracia, J.-Y. Delétage, and H. Frémont, "Effect of thermal and vibrational combined ageing on QFN terminal pads solder reliability," *Microelectronics Reliability*, vol. 114, p. 113883, 2020.

[4] C. Cui, F. F. Chamasemani, P. Paulachan, *et al.*, "Correlative, ML-based and non-destructive 3D analysis of intergranular fatigue cracking in SAC305-Bi solder balls," *npj Materials Degradation*, vol. 8, no. 40, 2024.

[5] K. Rathod, S. Desapogu, A. Jansche, T. Bernthaler, G. Schneider, D. Braun, and S. Diez, "Nondestructive defect detection in 3D X-ray microscopy data of ball grid array solder for void detection in solder joints using deep learning," *EDFA Tech. Art.*, vol. 26, no. 3, pp. 4–11, Aug. 2024, doi: 10.31399/asm.edfa.2024-3.p004.

[6] V. Thukral, M. van Soestbergen, J. J. M. Zaal, R. Roucou, R. T. H. Rongen, W. D. van Driel, and G. Q. Zhang, "Board level vibration test method of components for automotive electronics: State-of-the-art approaches and challenges," *Microelectronics Reliability*, vol. 139, p. 113563, 2022.

[7] I. Gershman and J. B. Bernstein, "Structural health monitoring of solder joints in QFN package," *Microelectronics Reliability*, vol. 52, no. 12, pp. 3011–3016, 2012.

[8] S. Bi, C. Wang, B. Wu, *et al.*, "A comprehensive survey on applications of AI technologies to failure analysis of industrial systems," *Engineering Failure Analysis*, vol. 148, p. 107172, 2023.

[9] S. Morin, V. Dubourg, and D. Nudzinski, "Reliability analysis and condition monitoring of SAC+ solder joints under high thermomechanical stress conditions using neural networks," in *Proc. IEEE Workshop on Microelectronics and Reliability*, 2023.

[10] X. Jiang, L. Yang, Q. Liu, Q. Peng, L. Cheng, and G. Li, "Comparison of fatigue life prediction methods for solder joints under random vibration loading," *Microelectronics Reliability*, 2019, doi: 10.1016/j.microrel.2019.02.008.

[11] T.-Y. Hung, C.-J. Huang, C.-C. Lee, C.-C. Wang, K.-C. Lu, and K.-N. Chiang, "Investigation of solder crack behavior and fatigue life of the power module on different thermal cycling period," *Microelectron. Eng.*, vol. 107, pp. 125–129, 2013, doi: 10.1016/j.mee.2012.09.014.

[12] R. Khanam and M. Hussain, "YOLOv11: An overview of the key architectural enhancements," *arXiv preprint* arXiv:2410.17725, 2024. [Online]. Available: https://arxiv.org/abs/2410.17725

[13] Y. Zhou, C. Guo, X. Wang, Y. Chang, and Y. Wu, "A survey on data augmentation in the large-model era," *arXiv preprint* arXiv:2401.15422, 2024. [Online]. Available: https://arxiv.org/abs/2401.15422

[14] W. Li, Z. Fang, D. Song, W. Nie, X. Li, and A.-A. Liu, "Multi-view contrastive learning for unsupervised 3D model retrieval and classification," *Signal Process. Image Commun.*, vol. 138, p. 117333, 2025, doi: 10.1016/j.image.2025.117333.

[15] S. Luan, C. Wang, H. Zhang, X. Gao, and Y. Mao, "DSCUNet: A multi-view adaptive fusion network for pavement crack segmentation," in *Proc. 7th Int. Conf. Electron. Inf. Technol. Comput. Eng. (EITCE)*, Xiamen, China, 2023, pp. 1672–1676, doi: 10.1145/3650400.3650679.

[16] H. Cui, X. Liu, and N. Huang, "Pulmonary vessel segmentation based on orthogonal fused U-Net++ of chest CT images," in *Proc. MICCAI*, Lecture Notes in Computer Science, vol. 11769, Springer, Cham, 2019, doi: 10.1007/978-3-030-32226-7-33.

[17] G. Piantadosi, M. Sansone, R. Fusco, and C. Sansone, "Multi-planar 3D breast segmentation in MRI via deep convolutional neural networks," *Artif. Intell. Med.*, vol. 103, 2020, doi: 10.1016/j.artmed.2019.101781.

[18] N. Audebert, A. Boulch, B. Le Saux, and S. Lefèvre, "Distance transform regression for spatially-aware deep semantic segmentation," *Comput. Vis. Image Underst.*, vol. 189, 2019, doi: 10.1016/j.cviu.2019.102809.

[19] J. Tian, W. Cheung, N. Glaser, Y.-C. Liu, and Z. Kira, "UNO: Uncertainty-aware noisy-OR multimodal fusion for unanticipated input degradation," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Paris, France, 2020, pp. 5716–5723, doi: 10.1109/ICRA40945.2020.9197266.

[20] S. Liu, J. Zha, J. Sun, Z. Li, and G. Wang, "EdgeYOLO: An edge-real-time object detector," *arXiv preprint* arXiv:2302.07483, Feb. 2023. [Online]. Available: https://arxiv.org/abs/2302.07483.

[21] C. Zhou, X. Li, C. C. Loy, and B. Dai, "EdgeSAM: Prompt-in-the-loop distillation for on-device deployment of SAM," *arXiv preprint* arXiv:2312.06660, Dec. 2023. [Online]. Available: https://arxiv.org/abs/2312.06660.

[22] Y. Wang, J. He, and S. Yu, "Crack-EdgeSAM: Self-prompting crack segmentation system for edge devices," *arXiv preprint* arXiv:2412.07205, Dec. 2024. [Online]. Available: https://arxiv.org/abs/2412.07205.