Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira, Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria
© 2026 ESREL 2026 Organizers. *Published by* Research Publishing, Singapore.
doi: 10.3850/ESREL2026061419_esrel26-p26093-cd

Enhancing PCB Reliability with Autoencoder-Based Anomaly Detection for Oxidation Degradation in Selective Soldering Nozzles

Josef Fleischmann

Institute for Digital Production Management, University of Applied Sciences, Ansbach, Germany. E-mail: j.fleischmann@hs-ansbach.de

Ana I. Sánchez

Department of Statistics and Operational Research, MEDASEGI research group, Universitat Politècnica de València, Valencia, Spain. E-mail: aisanche@eio.upv.es

Oscar Trull

Department of Statistics and Operational Research, Universitat Politècnica de València, Valencia, Spain. E-mail: otrull@eio.upv.es

Juergen Goehringer

Institute for Digital Production Management, University of Applied Sciences, Ansbach, Germany. E-mail: juergen.goehringer@hs-ansbach.de

According to the latest reports, selective soldering plays a critical role in the manufacturing of printed circuit boards (PCBs), especially for through-hole technology (THT) components, where molten solder is dispensed from a nozzle to create a continuous 360° wave around the part. The consistency of this wave is crucial for reliable electrical and mechanical connections. However, repeated exposure to high temperatures (up to 300°C) and molten solder (e.g., Sn-Pb or SAC alloys) causes nozzle oxidation, leading to impurity buildup and irregular waves. This degradation results in defective solder joints, compromising product quality, increasing rework costs, and potentially causing production downtime. Current monitoring relies on manual visual inspections, which are subjective and time-consuming, or rule-based image algorithms that are prone to noise and require frequent recalibration. To overcome these limitations, this paper investigates autoencoders—a type of neural network for unsupervised anomaly detection—that reconstruct input images and flag deviations via reconstruction errors. Data was collected from camera-monitored nozzles under controlled conditions. Preprocessing included grayscale conversion, Gaussian blurring, normalization, and resizing. A Flat (fully connected baseline) was evaluated. Models trained on operable images using mean squared error (MSE) loss; performance assessed with MSE distributions and receiver operating characteristic (ROC) curves, using area under the curve (AUC) as the primary metric. Results show the flat autoencoder achieving ~96% AUC with a MSE separation (oxidized > 0.02). This is the first application of autoencoders to real-time nozzle oxidation detection in selective soldering, filling a gap in AI for tool-specific degradation. It enables predictive maintenance, reducing defects.

Keywords: Autoencoders, anomaly detection, selective soldering, nozzle degradation, condition monitoring.

1. Introduction

In recent years, the application of machine learning algorithms to anomaly detection in industrial contexts has increased significantly (Givnan et al. 2022, Maggipinto et al. 2022, Neloy & Turgeon 2024, Kyeong et al. 2023). These approaches facilitate the automated identification of deviations from standard operating conditions in complex systems, where conventional techniques, such as rule-based inspection, are frequently ineffective. Consequently, machine learning-based anomaly detection has become a key enabling technology in predictive maintenance across multiple industrial sectors.

Selective soldering is a critical process in the manufacturing of printed circuit boards (PCBs), especially for through-hole technology (THT) components. In selective soldering, components are joined by immersion in molten solder delivered through a nozzle that produces a continuous 360° solder wave around the lead (Diepstraten 2016). The uniformity and stability of this wave are essential for achieving reliable electrical connections and sufficient mechanical integrity (Qasaimeh et al. 2024). However, as the soldering nozzle is exposed to repeated heating (up to 300°C) and solder oxidation, its surface gradually accumulates impurities, resulting in irregular and unstable soldering waves (McMaster et al. 2023). Oxidation occurs through repeated thermal cycles and exposure to molten solder, resulting in a surface buildup similar to rust via chemical reactions such as $2Sn + O_2 \rightarrow 2SnO$. This degradation can lead to defective solder joints, negatively impacting product quality and reliability (Seehase et al. 2018).

Different strategies have been proposed to monitor the condition of soldering nozzles. In current industrial practice, periodic visual inspections by operators remain common, although they are subjective, non-continuous, and typically capable of detecting oxidation only at advanced stages (Schmidt et al., 2021). Alternatively, rule-based algorithms have been employed to estimate oxidation levels from images captured by cameras monitoring the nozzle, converting them to binary representations to quantify the degree of oxidation (Zhang et al. 2023). Although partially effective, these algorithms are prone to false readings due to background noise and require frequent recalibration (Eisentraut et al. 2025). Figure 1 shows an example of selective soldering of a THT component.

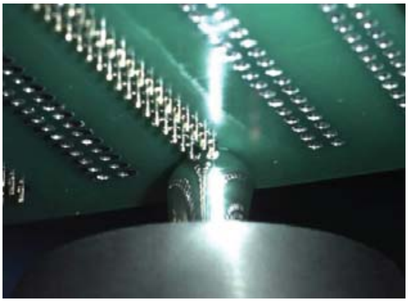

Fig. 1. Selective soldering of a THT component.

In this context, this paper studies the use of autoencoders for anomaly detection in soldering nozzles. Autoencoders are unsupervised neural network models that learn compact representations of normal operating conditions and identify anomalies by analyzing reconstruction errors (Presenti et al. 2022). Their successful application in predictive maintenance, condition monitoring, and semiconductor manufacturing highlights their suitability for selective soldering processes (Mendizabal-Arrieta et al. 2025, Rakshitha et al. 2025). In the proposed approach, autoencoders are trained exclusively on images of operable, non-oxidized nozzles. Deviations from the learned normal behavior are subsequently used to automatically identify oxidation and surface degradation, thereby enabling early maintenance actions and reducing the probability of defective solder joints.

The results obtained show that the proposed method is capable of reliably detecting nozzle degradation from image data, even under varying background and illumination conditions. These findings contribute to the development of automated, data-driven monitoring systems for selective soldering processes.

This paper is structured as follows. Section 2 presents the basic concepts related to autoencoders and motivates the architectural choice. Section 3 describes the application case and the results obtained. Section 4 discusses the findings in the context of related work. Finally, Sections 5 and 6 present the conclusions and outlook.

2. Autoencoders

Autoencoders (AEs) (Xu et al. 2025 & McMaster et al. 2023), are neural network architectures designed for unsupervised learning. Their main objective is to learn a compact latent representation of the input data while minimizing the reconstruction error between the original and reconstructed signal. To achieve this, the network projects the input into a lower-dimensional latent space that captures the underlying structure and correlations in the data and subsequently reconstructs the original signal from this representation. When the model is trained exclusively on data corresponding to normal operating conditions, samples that deviate from this behavior cannot be accurately reconstructed, resulting in significantly higher reconstruction errors. This characteristic makes autoencoders particularly appropriate for anomaly detection in complex industrial systems. In the context of selective soldering, for instance, nozzle oxidation or surface degradation can lead to anomalous image patterns that manifest as an increased reconstruction error.

An autoencoder consists of two components, an encoder and a decoder, as illustrated in Figure 2. The encoder learns a nonlinear mapping from the input space $\boldsymbol{x} \in \Re^n$ to a latent space $\boldsymbol{z} \in \Re^m$, being $m < n$, defined as:

$$\boldsymbol{z} = f(\boldsymbol{x}) \tag{1}$$

The decoder performs the inverse mapping, reconstructing the input from the latent representation:

$$\hat{\boldsymbol{x}} = g(\boldsymbol{z}) \tag{2}$$

that reconstructs $\hat{\boldsymbol{x}}$.

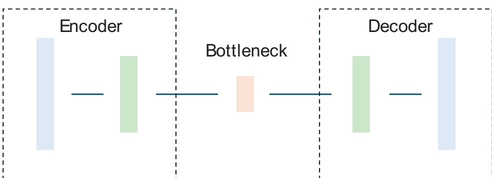

Fig. 2. Basic autoencoder scheme.

The autoencoder is trained by minimizing the Mean Squared Error (MSE) between input and reconstruction, and the resulting reconstruction error serves as an anomaly score. The detection threshold is determined using the Youden index on the Receiver Operating Characteristic (ROC) curve, the point that maximizes the sum of True-Positive-Rate (TPR) and specificity (1 – False-Positive-Rate (FPR)), thereby achieving the best overall balance between detection performance and false alarms (Givnan et al. 2022). This approach ensures reliable detection of oxidized nozzles while effectively controlling the trade-off between true positives and false positives.

To improve robustness, images are pre-processed through grayscale conversion, Gaussian blurring, normalization to the range [0, 1], and resizing to 256 × 256 pixels, thereby reducing noise and ensuring consistent input for accurate reconstruction.

In the literature, various autoencoder variants have been proposed for anomaly detection, including convolutional autoencoders (CAEs) that preserve spatial structure through convolutional and pooling layers (Lu et al. 2024, Maggipinto et al. 2022), variational autoencoders (VAEs) with probabilistic latent spaces, and vision transformer-based architectures (Nguyen et al. 2024). While CAEs have demonstrated strong performance for spatially localized defects such as scratches or dents (Bergmann et al. 2021), the anomaly detection task addressed in this work differs fundamentally: nozzle oxidation manifests as a global surface change affecting the entire nozzle tip, rather than a spatially confined defect. A fully connected (flat) autoencoder is therefore employed, as it operates on the holistic image representation and is well-suited to capture image-wide distributional shifts associated with oxidation. Furthermore, its low parameter count and fast inference time make it particularly suitable for real-time deployment in production environments operating at 30 fps. Recent comparative studies have shown that more complex architectures do not consistently outperform simpler ones, particularly when training data is limited and hyperparameter tuning is constrained (Nguyen et al. 2024, Neloy & Turgeon 2024).

The flat autoencoder consists of fully connected feed-forward layers with a single hidden layer in both the encoder and the decoder. Given an input image $x \in \Re^n$, flattened into a one-dimensional vector, the encoder learns a mapping to a latent representation $z \in \Re^m$, defined as:

$$\boldsymbol{z} = \phi(\boldsymbol{W}_e \boldsymbol{x} + \boldsymbol{b}_e) \tag{3}$$

where $\phi(\cdot)$ is a non-linear activation function, and $\boldsymbol{W}_e$ and $\boldsymbol{b}_e$ denote the encoder weights and biases, respectively. The decoder reconstructs the input according to:

$$\hat{\boldsymbol{x}} = \phi(\boldsymbol{W}_d \boldsymbol{z} + \boldsymbol{b}_d) \tag{4}$$

3. Application case

The application case is focused on the practical application of autoencoders for anomaly detection in selective soldering nozzles using real-world nozzles images to identify oxidation anomalies. A high-resolution camera with a resolution of 1024x1024 pixels and a frame rate of 30 frames per second was installed to acquire images of the nozzle surface under representative operating conditions. The camera was mounted at a fixed distance on the side of the soldering nozzle to ensure a consistent field of view across all acquisitions. Images were captured during test production cycles under varying but representative illumination conditions inherent to the production environment. No additional or controlled lighting was applied, meaning the dataset reflects real-world variability in background and illumination. Each image was cropped to a region of interest centered on the nozzle tip before preprocessing.

Image labeling was performed by experienced machine operators with domain expertise in selective soldering processes. Each image was visually assessed and assigned to one of two categories: operable (no visible oxidation or degradation affecting nozzle function) or oxidized (visible surface discoloration, buildup, or irregular texture indicative of oxidation). In cases of ambiguity, a second operator reviewed the label to ensure consistency. No intermediate degradation stages were defined, as the primary objective was binary anomaly detection.

Visually, operable nozzles exhibit a uniform, metallic surface with consistent reflective properties across the nozzle tip. Oxidized nozzles, in contrast, display a range of degradation patterns, including irregular texture caused by oxide buildup and in advanced stages, visible pitting or flaking of the surface layer. The degree of oxidation varies considerably within the dataset, from early-stage discoloration barely distinguishable from normal surface variation to severely degraded nozzles with pronounced structural changes. This variability, combined with the uncontrolled illumination conditions of the production environment, contributes to the challenge of the detection task and accounts for the observed false negatives at the current detection threshold.

A total of 1371 images of operable nozzles were used for training the autoencoder, with a 10% validation split (resulting in approximately 1234 images for training and 137 images for validation). For testing, 1640 images were used, comprising 264 operable (OK) nozzles and 1376 oxidized (fail) nozzles. The training set consists exclusively of operable nozzle images, as the autoencoder is trained in an unsupervised manner to learn the representation of normal operating conditions. The test set was intentionally composed with a higher proportion of oxidized images (1376 oxidized vs. 264 operable) to reflect the practical scenario in which degradation detection is the primary objective and to ensure statistically meaningful evaluation of detection performance across both classes.

The autoencoder model was trained using the MSE loss function, which measures pixel-level discrepancies between original and reconstructed images.

Key hyperparameters were optimized through iterative experimentation, including the learning rate (set to 0.001 using the Adam optimizer), batch size (32), and number of epochs (24). These values were selected based on the minimization of validation loss and reconstruction quality. Performance was assessed by analyzing MSE distributions for operable and oxidized nozzles, and by plotting ROC curves, with the Area Under the Curve (AUC) serving as the primary metric.

Figure 3 illustrates sample reconstructions produced by flat autoencoder architecture, demonstrating a high fidelity in reconstructing operable nozzles, while the reconstructions of oxidized nozzles showed significant degradation.

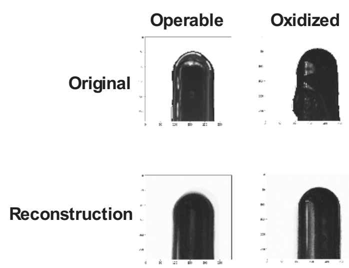

Fig. 3. Autoencoder Reconstructions of Operable and Oxidized Nozzles.

Additionally, Figure 4 displays the MSE distribution for operable and oxidized nozzles. The distribution reveals a separation between the two classes: operable nozzles exhibit low MSE values (typically below 0.02, while oxidized nozzles show higher MSE values (often above 0.02, with a broader spread, reflecting significant reconstruction errors due to degradation anomalies. This bimodal pattern supports effective threshold-based anomaly detection.

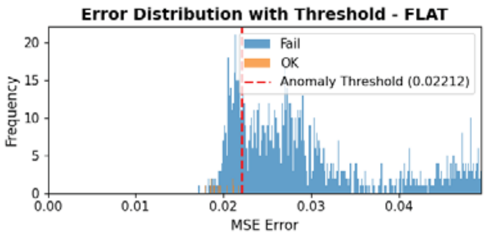

Fig. 4. MSE Distributions for Operable and Oxidized Nozzles (blue=oxidized, orange=operable).

The model performance is evaluated using the ROC curve which is presented in Figure 5 where it is observed that the AUC is 96%, indicating solid classification performance and minimal false positives. Visual inspection confirmed that these architectures produced high-fidelity reconstructions for operable nozzles, while significantly degrading those for oxidized ones.

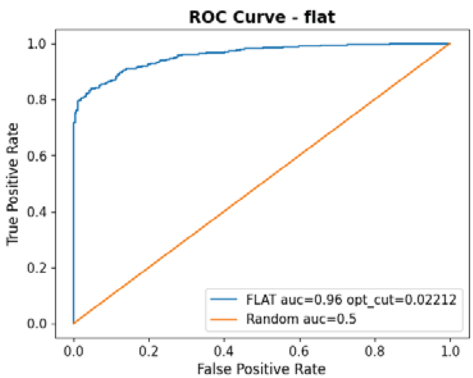

Fig. 5. ROC curve for flat autoencoder.

Figure 6 presents the confusion matrix for the flat autoencoder, evaluated on the full independent test set of 1640 images (264 operable/OK and 1376 oxidized/fail nozzles) using the optimal MSE threshold derived from the ROC curve.

It shows strong performance with 252 true negatives (correctly classified operable nozzles), 1153 true positives (correctly identified oxidized nozzles), 12 false positives, and 223 false negatives, resulting in an overall accuracy of 85.67%, precision of 98.97%, recall of 83.79%, and F1 score of 90.75%. This indicates very high precision in flagging anomalies (very few false alarms) while achieving good sensitivity in detecting the majority of degraded nozzles, with opportunities to further reduce false negatives through architectural refinements or threshold optimization.

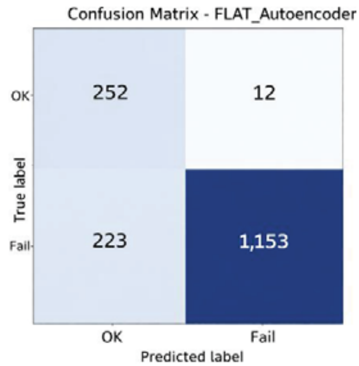

Fig. 6. Confusion matrix for the flat autoencoder.

4. Discussion

The flat autoencoder achieves an AUC of 96% and a precision of 98.97% on the independent test set, demonstrating that even a comparatively simple architecture can provide reliable anomaly detection for nozzle oxidation monitoring. The high precision indicates that nearly all flagged nozzles are genuinely degraded, which is critical in production environments where false alarms disrupt workflows.

The choice of a fully connected architecture over convolutional alternatives is motivated by the nature of the anomaly: oxidation causes a diffuse, global surface change rather than a localized defect. Convolutional autoencoders, which excel at detecting spatially confined anomalies through local receptive fields (Bergmann et al. 2021), offer limited advantage when the relevant signal is distributed across the entire image. In the MVTec AD benchmark, CAE-based methods achieve image-level AUROC scores ranging from 73% to 87% depending on the object category, with performance varying significantly across defect types (Bergmann et al. 2021). This underscores that architectural complexity does not guarantee superior performance, particularly for problems where global image features dominate.

A practical consideration further supports this architectural choice: the flat autoencoder's low parameter count enables inference times well below the 33 ms threshold required for real-time monitoring at 30 fps, facilitating deployment on standard industrial hardware without GPU acceleration. This is consistent with findings by Neloy and Turgeon (2024), who observe that simpler autoencoder configurations can match or exceed the performance of deeper architectures when the feature space is well-suited to the detection task.

The observed recall of 83.79% indicates that approximately 16% of oxidized nozzles are not detected at the current threshold. This is primarily attributable to early-stage oxidation, where surface changes are subtle and reconstruction errors remain below the detection threshold. Adjusting the threshold would increase recall at the cost of precision, representing a trade-off that must be calibrated to the specific production requirements. Future work may address this limitation through multi-scale feature extraction or ensemble approaches.

An empirical comparison with alternative architectures on the present dataset is planned as part of ongoing work; however, the literature evidence presented above provides a principled justification for the architectural choice in this study.

5. Conclusion

This paper presented an autoencoder-based approach for automated oxidation detection in selective soldering nozzles. Using a fully connected autoencoder trained exclusively on images of operable nozzles, the method achieves an AUC of 96%, a precision of 98.97%, and an F1 score of 90.75% on an independent test set of 1640 images. The architectural choice was motivated by the global nature of oxidation anomalies and the requirement for real-time inference in production environments, and is supported by comparative evidence from the anomaly detection literature indicating that simpler architectures can be effective when matched to the problem structure. The clear separation of MSE distributions between operable and oxidized nozzles supports threshold-based classification for reliable deployment.

In summary, autoencoder-based anomaly detection provides an efficient solution for monitoring soldering nozzle degradation.

6. Outlook

Future research will investigate threshold optimization strategies to improve recall for early-stage oxidation while maintaining high precision. Additionally, the integration of temporal information from consecutive frames may enable trend-based detection of gradual nozzle degradation. Expanding the dataset to include intermediate degradation stages and multiple nozzle geometries will further strengthen the generalizability of the approach.

References

Diepstraten, G. (2016). Design Improvements for Selective Soldering Assemblies. *IPC APEX Expo*.

Bergmann, P., M. Fauser, D. Sattlegger, and C. Steger (2021). The MVTec Anomaly Detection Dataset: A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. International Journal of Computer Vision 129(4), 1038–1059.

Eisentraut, L., J. Hosch, M. Roytenberg, A. Benecke, P. Penava, and R. Buettner (2025). Defect Detection in Industrial Soldering Processes Using Machine Learning: A Critical Literature Review. *IEEE Access,13,* 10909447.

Givnan, S., C. Chalmers, P. Fergus, S. Ortega-Martorell, and T. Whalley (2022). Anomaly Detection Using Autoencoder Reconstruction upon Industrial Motors. *Sensors 22(9)*, 3166.

Kyeong, Y., D. Kim, E. Kim, S. Park , Y. Jin, J. Kim, and S. Woo (2023, June). P-167: Autoencoder with Variables Added Model for Manufacturing Process Anomaly Detection Using Multivariate Time-series Data that Minimizes Data Information Loss. In *SID Symposium Digest of Technical Papers* (Vol. 54, No. 1, pp. 1587-1590).

Lu, S., W. Zhang, H. Zhao, H. Liu, N. Wang, and H. Li (2024). Anomaly detection for medical images using heterogeneous auto-encoder. *IEEE Transactions on Image Processing*, *33*, 2770-2782.

Maggipinto, M., A. Beghi, and G.A. Susto (2022). A deep convolutional autoencoder-based approach for anomaly detection with industrial, non-images, 2-dimensional data: A semiconductor manufacturing case study. *IEEE Transactions on Automation Science and Engineering*, *19*(3), 1477-1490.

McMaster, S., A. Cobley, N. Monk, and J. Graves (2023, November). Selective soldering nozzles: insights into wear mechanisms and future developments. In *IPC APEX Technical Conference*.

McMaster, S., K. Witham, M. Stephenson, E. Glover, and E. Groves (2024, May). Process Challenges for Selective Soldering: Examining parameters for optimal processing. In *ECWC16 at IPC APEX EXPO 2024*. IPC International.

Nguyen, H.H., C.N. Nguyen, X.T. Dao, Q.T. Duong, D.P.T. Kim, and M.-T. Pham (2024). Variational Autoencoder for Anomaly Detection: A Comparative Study. arXiv preprint arXiv:2408.13561.

Mendizabal-Arrieta, I., H. Álvarez, D. Aguinaga, A. Pellicer, J.R. Sánchez, and F. Torres (2025). Deep learning-based solder joint defect detector. *The International Journal of Advanced Manufacturing Technology*, *137(9)*, 5133-5147.

Neloy, A. and M. Turgeon (2024). A comprehensive study of auto-encoders for anomaly detection: Efficiency and trade-offs. *Machine Learning with Applications*, *17*, 100572.

Presenti, A., Z. Liang, L.F. Alves Pereira, J. Sijbers, and J. De Beenhouwer (2022). Automatic anomaly detection from X-ray images based on autoencoders. *Nondestructive Testing and Evaluation*, *37*(5), 552-565.

Qasaimeh, Q., H. Li, S.D. Hamasha, J. Evans, and J. Liu (2024). Interpretable Machine Learning Models Can Outperform Statistical Models in Solder Joint Reliability. In *2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)* (1-9). IEEE.

Rakshitha, G., J. Sinchana, S. Manuj, K.N.Pranathi, and Anithadb (2025). Solder Joint Defect Detection in PCB Chip Components Using Machine Learning. *International Journal of Research Publication and Reviews*, *6 (8)*, 51530.

Seehase, D., C. Kohlen, A. Neiser, A. Novikov, and M. Nowottnick (2018). Selective soldering on printed circuit boards with endogenous induction heat at appropriate susceptors. *Periodica Polytechnica Electrical Engineering and Computer Science*, *62*(4), 172-180.

Schmidt, K., D. Rauchensteiner, C. Voigt, N. Thielen, J. Bönig, G. Beitinger and J. Franke (2021, June). An automated optical inspection system for PIP solder joint classification using convolutional neural networks. In *2021 IEEE 71st Electronic Components and Technology Conference (ECTC)* (2205-2210). IEEE.

Xu, Y., H. Wu, Y. Liu, and X. Zhang (2025). PCB Electronic Component Soldering Defect Detection Using YOLO11 Improved by Retention Block and Neck Structure. *Sensors*, *25*(11), 3550.

Zhang, Z., W. Zhang, D. Zhu, Y. Xu, and C. Zhou (2023). Printed circuit board solder joint quality inspection based on lightweight classification network. *IET Cyber-Systems and Robotics*, *5*(4), e12101.