Optimizing Deep Learning Performance in PCB X-ray Inspection through Synthetic Data Tuner

Shajib Ghosh, Md Mahfuz Al Hasan, Patrick Craig, Antika Roy, Nitin Varshney, Sanjeev J. Koppal, Hamed Dalir, and Navid Asadizanjani

Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA

ABSTRACT

In the domain of printed circuit board (PCB) X-ray inspection, the effectiveness of deep learning models greatly depends on the availability and quality of annotated data. The utilization of data augmentation techniques, particularly through the utilization of synthetic data, has emerged as a promising strategy to improve model performance and alleviate the burden of manual annotation. However, a significant question remains unanswered: What is the optimal amount of synthetic data required to effectively augment the dataset and enhance model performance? This study introduces the Synthetic Data Tuner, a comprehensive framework developed to address this crucial question and optimize the performance of deep learning models for PCB X-ray inspection tasks. By employing a combination of cutting-edge deep learning architectures and advanced data augmentation techniques, such as generative adversarial networks (GANs) and variational autoencoders (VAEs), the Synthetic Data Tuner systematically assesses the impact of different levels of synthetic data integration on model accuracy, robustness, and generalization. Through extensive experimentation and rigorous evaluation procedures, our results illustrate the intricate relationship between the quantity of synthetic data and model performance. We elucidate the phenomenon of diminishing returns, where model performance reaches a saturation point beyond a specific threshold of synthetic data augmentation. Moreover, we determine the optimal balance between synthetic and real data, achieving a harmonious equilibrium that maximizes performance improvements while mitigating the risk of overfitting. Additionally, our findings emphasize the significance of data diversity and quality in the generation of synthetic data, highlighting the importance of domain-specific knowledge and context-aware augmentation techniques. By providing insights into the complex interplay between synthetic data augmentation and deep learning model performance, the Synthetic Data Tuner not only advances the current state-of-the-art in PCB X-ray inspection but also offers valuable insights and methodologies applicable to various computer vision and industrial inspection domains.

Keywords: Synthetic Data Tuner, PCB X-ray Inspection, Deep Learning, Data Augmentation, Performance Optimization

1. INTRODUCTION

1.1 Motivation

The significance of PCB X-ray inspection in ensuring the reliability and functionality of electronic devices is crucial, as it plays a pivotal role in detecting defects that are often imperceptible to the naked eye.[1–4] Nonetheless, a major challenge in this field is the procurement of annotated data, which is both time-consuming and labor-intensive. The process of manual annotation necessitates expert knowledge and meticulous effort, rendering it a costly undertaking. The scarcity of annotated data poses a hindrance to the efficient training of deep learning models, which rely on extensive datasets to achieve high accuracy and robustness.[5–7] To surmount this barrier, the utilization of synthetic data has emerged as a potent solution in the realm of deep learning.[8] Synthetic data, produced through sophisticated methods such as generative adversarial networks (GANs) and variational autoencoders (VAEs), provides the benefit of generating diverse and abundant datasets without the requirement

Further author information: (Send correspondence to S.G. or N.A.)
S.G.: E-mail: shajib.ghosh@ufl.edu; N.A.: E-mail: nasadi@ece.ufl.edu

Applications of Machine Learning 2024, edited by Michael Zelinski, Tarek Taha, Barath Narayanan, Abdul Awwal, Khan Iftekharuddin, Proc. of SPIE Vol. 13138, 131380N

for extensive manual annotation.[9] Through the incorporation of synthetic data into the training regimen, models can encounter a wider range of variations, thereby enhancing their capacity to generalize and perform effectively on real-world data.[8] This strategy not only alleviates the challenges associated with data annotation but also expedites the development cycle of deep learning applications in PCB X-ray inspection, establishing it as a crucial domain for research and innovation.

1.2 Problem Statement

The importance of having an appropriate quantity of synthetic data in PCB X-ray inspection is emphasized by the need to improve the performance of deep learning models without facing the high costs linked to manual annotation. Striking a harmonious blend between synthetic and authentic data proves to be a notable challenge, given that an excessive reliance on synthetic data may result in overfitting and diminished model generalization, whereas an insufficient supply of synthetic data hinders the ability to cultivate diverse training models.[9] This intricate balance holds paramount significance, as it dictates the model's capacity to effectively identify and categorize anomalies in real-life scenarios, thus demanding a methodical approach to pinpoint the ideal proportion of synthetic versus authentic data. The creation of frameworks that can dynamically adjust this equilibrium is crucial for progressing the efficiency of PCB X-ray inspection systems and ensuring their relevance in various and fluctuating operational settings.

1.3 Objectives

The Synthetic Data Tuner framework is being introduced in this study with the aim of systematically evaluating the impact of synthetic data on deep learning models utilized in PCB X-ray inspection. The main objectives include determining the optimal balance between synthetic and real data, as well as improving model performance in terms of accuracy, robustness, and generalization. By utilizing advanced data augmentation techniques like GANs and VAEs, the framework assesses how different levels of synthetic data integration affect model outcomes. This research aims to offer valuable insights into the most effective strategies for leveraging synthetic data to enhance model performance in industrial inspection applications through thorough experimentation and analysis.

2. RELATED WORKS

2.1 Data Augmentation Techniques

Traditional data augmentation approaches are crucial for enhancing deep learning algorithms by artificially increasing the size and variety of training datasets, thereby improving generalization and reducing overfitting. Common augmentation techniques include geometric transformations, color space augmentations, and noise injection. Geometric transformations, such as rotation, flipping, scaling, random cropping, and translation, help models recognize components from various orientations and simulate different feature sizes, enhancing robustness to variations. Color space transformations adjust RGB pixel values to address illumination changes and color biases, with simple operations like brightness adjustment yielding significant improvements. Noise injection introduces random noise to images, boosting the resilience of convolutional neural networks to real-world data imperfections.[8] In deep learning-based augmentations, feature space augmentations leverage neural networks to map high-dimensional inputs into lower-dimensional representations for data manipulation, with autoencoders being particularly effective. Generative modeling, including GANs and VAEs, produces synthetic data that maintains original dataset characteristics, enhancing training data quality and diversity.[8] Adversarial training employs two networks with opposing objectives to explore augmentation spaces, generating challenging augmented instances that improve model robustness. These techniques collectively enhance model performance by providing diverse and high-quality training data.

2.2 Generative Models for Data Synthesis

Generative models, specifically GANs and VAEs, have transformed the landscape of synthetic data generation within the realm of deep learning. GANs are comprised of a pair of neural networks, a generator and a discriminator, which undergo concurrent training via adversarial mechanisms. The generator fabricates synthetic data, while the discriminator assesses their genuineness, ultimately resulting in the production of exceptionally realistic data.[9] On the contrary, VAEs employ a probabilistic approach, converting input data into a latent space prior to reconstructing it in order to generate novel, similar data. This particular technique guarantees that the produced data is both varied and a faithful representation of the original dataset.[10] The utilization of both GANs and VAEs has showcased substantial promise in enriching datasets, consequently amplifying model efficacy across a multitude of domains, including PCB X-ray scrutiny.

3. METHODOLOGY

3.1 Synthetic Data Tuner Framework

The Synthetic Data Tuner framework (as illustrated in Figure 1) is a comprehensive system created to enhance the performance of deep learning models in PCB X-ray inspection by utilizing synthetic data strategically. The framework consists of multiple essential components that work together seamlessly towards this objective. The data preprocessing module manages the initial processing of raw PCB images and labels, ensuring proper resizing, normalization, and formatting for further processing. The VAE and GAN training modules focus on the development of VAEs and GANs to generate high-quality synthetic images that supplement the original dataset. The synthetic data generation module utilizes the trained VAE and GAN models to produce synthetic data, which is then merged with real data to form a robust training dataset. Lastly, the training and evaluation module utilizes a ResNet model to systematically evaluate the impact of different ratios of synthetic to real data, employing stratified k-fold cross-validation to assess model performance across various metrics. This holistic approach enables precise adjustment of synthetic data levels, resulting in improved model accuracy, resilience, and generalization in PCB X-ray inspection tasks.

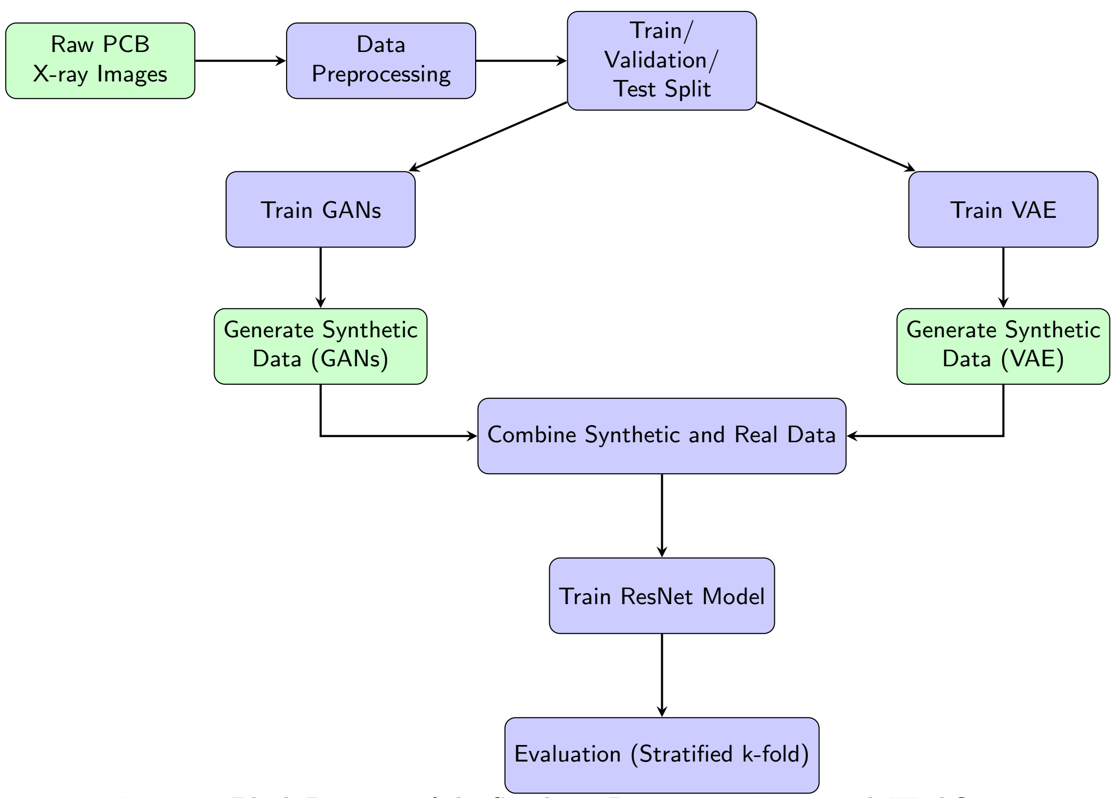

Figure 1: Block Diagram of the Synthetic Data Tuner Framework Workflow

3.2 Data Preprocessing

The data preprocessing pipeline for the Synthetic Data Tuner framework encompasses a sequence of meticulous procedures to ready raw PCB images and their corresponding labels for model training. Initially, the images undergo resizing, conversion to RGB format, and normalization utilizing the PyTorch[11] and torchvision libraries[12] to ensure consistency in input dimensions and pixel intensity values. The PCBImageDataset class systematically uploads the images and labels, making use of the PIL library[13] for image manipulation and PyTorch[11] for tensor operations. Techniques for data augmentation, such as random horizontal and vertical flips, as well as random rotations, are employed to boost dataset diversity and resilience. A custom collate function guarantees that segmentation labels of different lengths are suitably padded, thereby enabling efficient batch processing during model training. This comprehensive preprocessing routine plays a crucial role in enhancing the performance and generalization capabilities of deep learning models in PCB X-ray inspection.

3.3 Model Training

The Synthetic Data Tuner framework utilizes a blend of Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and ResNet architectures to boost model performance in categorizing components on PCB X-ray images. Different GAN models, such as DCGAN, WGAN-GP, and SNGAN, are used to produce synthetic images by amplifying noise for realistic depictions.[9] DCGAN employs a discriminator to differentiate between real and synthetic images,[14] WGAN-GP includes a gradient penalty for better training stability,[15,16] and SNGAN applies spectral normalization to improve training dynamics.[17] The VAE model, utilizing a ResNet encoder, condenses images into a latent space and then reconstructs them through a decoder to ensure varied and top-notch synthetic data.[10] The ResNet architecture[18] is utilized for the classification task, utilizing deep residual learning capabilities for precise PCB component classification.

3.4 Loss Functions and Evaluation Metrics

To evaluate the performance of the Synthetic Data Tuner framework, we employ various loss functions and evaluation metrics tailored to the specific models and tasks within the framework. This section details the mathematical formulation of these losses and metrics.

3.4.1 Loss Functions

For the GAN models, the loss functions are defined as follows:

1. **DCGAN Loss**: The objective of the DCGAN is to minimize the adversarial loss. The generator loss $\mathcal{L}_G$ and discriminator loss $\mathcal{L}_D$ are given by:

$$\mathcal{L}_D = -\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}[\log D(\mathbf{x})] - \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[\log(1 - D(G(\mathbf{z})))] \tag{1}$$

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[\log(D(G(\mathbf{z})))] \tag{2}$$

where D is the discriminator, G is the generator, $\mathbf{x}$ represents real data samples, and $\mathbf{z}$ represents the noise vector sampled from a prior distribution $p_{\mathbf{z}}$.

2. **WGAN-GP Loss**: The WGAN-GP modifies the standard GAN loss with a gradient penalty to ensure Lipschitz continuity.[19] The discriminator loss with the gradient penalty $\mathcal{L}_D$ and the generator loss $\mathcal{L}_G$ are formulated as:

$$\mathcal{L}_D = -\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}[D(\mathbf{x})] + \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[D(G(\mathbf{z}))] + \lambda\mathbb{E}_{\hat{\mathbf{x}}\sim p_{\hat{\mathbf{x}}}}[(\|\nabla_{\hat{\mathbf{x}}}D(\hat{\mathbf{x}})\|_2 - 1)^2] \tag{3}$$

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[D(G(\mathbf{z}))] \tag{4}$$

where λ is a regularization parameter, and $\hat{\mathbf{x}}$ is a random interpolation between real and generated data samples.

3. **SNGAN Loss**: The SNGAN uses a similar adversarial loss function to DCGAN but incorporates spectral normalization. The loss functions are:

$$\mathcal{L}_D = -\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}[\log D(\mathbf{x})] - \mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[\log(1 - D(G(\mathbf{z})))] \tag{5}$$

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z}\sim p_{\mathbf{z}}}[\log(D(G(\mathbf{z})))] \tag{6}$$

where spectral normalization stabilizes the training by controlling the Lipschitz constant of the discriminator.

For the VAE, the loss function combines the reconstruction loss and the Kullback-Leibler divergence:[20]

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})\|p_{\mathbf{z}}) \tag{7}$$

where $q_\phi(\mathbf{z}|\mathbf{x})$ is the encoder output, $p_\theta(\mathbf{x}|\mathbf{z})$ is the decoder output, and $p_\mathbf{z}$ is the prior distribution over the latent variable $\mathbf{z}$.

For the ResNet model used in component classification, the cross-entropy loss is utilized:

$$\mathcal{L}_{\text{CE}} = -\sum_{i=1}^{N} y_i \log(\hat{y}_i) \tag{8}$$

where y_i is the true label, $\hat{y}_i$ is the predicted probability for class i, and N is the total number of samples.

3.4.2 Evaluation Metrics

To evaluate the performance of the ResNet model in component classification, we employ the following metrics:

1. **Accuracy**:

$$\text{Accuracy} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}(\hat{y}_i = y_i) \tag{9}$$

where N is the total number of samples, $\hat{y}_i$ is the predicted label, and y_i is the true label.

2. **Precision**:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{10}$$

where TP is the number of true positives and FP is the number of false positives.

3. **Recall**:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{11}$$

where FN is the number of false negatives.

4. **F1 Score**:

$$\text{F1 Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{12}$$

These evaluation metrics provide a comprehensive assessment of the model's performance, capturing both its accuracy and its ability to correctly classify components on PCB X-ray images.

4. EXPERIMENTATION AND RESULTS

4.1 Experimental Setup

The experiments were conducted using an NVIDIA GeForce RTX 2080 Ti GPU, operating on a Red Hat Enterprise Linux (RHEL) Server version 7.9, equipped with an Intel® Core™ i9-10980XE Extreme Edition Processor clocked at a base frequency of 3.00 GHz. The dataset, collected with the facilities of the University of Florida and available upon request at https://physicaldb.ece.ufl.edu/index.php/segpcbx-dataset/, was split into training, validation, and testing sets using a stratified strategy to ensure balanced representation of component classes.

Data preprocessing included resizing images, normalizing pixel values, and augmenting the dataset with transformations like random flips and rotations. This preprocessing played a vital role in standardizing the input for models and boosting their resilience.

GANs (DCGAN, WGAN-GP, SNGAN) and VAEs were trained with different hyperparameters, such as various latent dimensions, learning rates, and batch sizes, to produce synthetic PCB images. Training of these generative models was directed by adversarial and reconstruction losses. The ResNet model, utilized for component classification, underwent training on a merged dataset of real and synthetic images, and its performance was assessed through cross-entropy loss.

Evaluation metrics involved accuracy, precision, recall, and F1 score, offering a thorough evaluation of the model's classification effectiveness. The training and evaluation procedures made use of stratified k-fold cross-validation to ensure robust and dependable outcomes. The entire project, including the codes and results, is available at the GitHub repository: https://github.com/shajibghosh/SyntheticDataTuner.git.

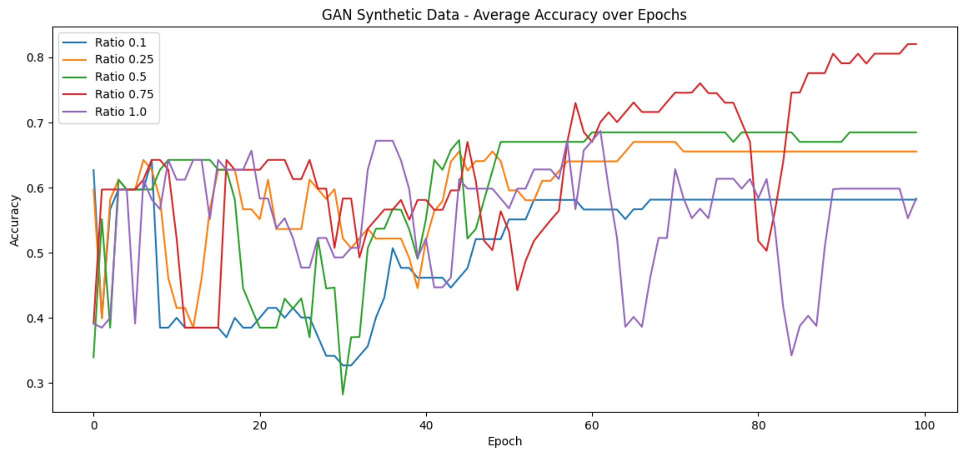

(a) Average accuracy vs. epochs plots of the classification model for different ratios of synthetic data produced using GAN training and real data.

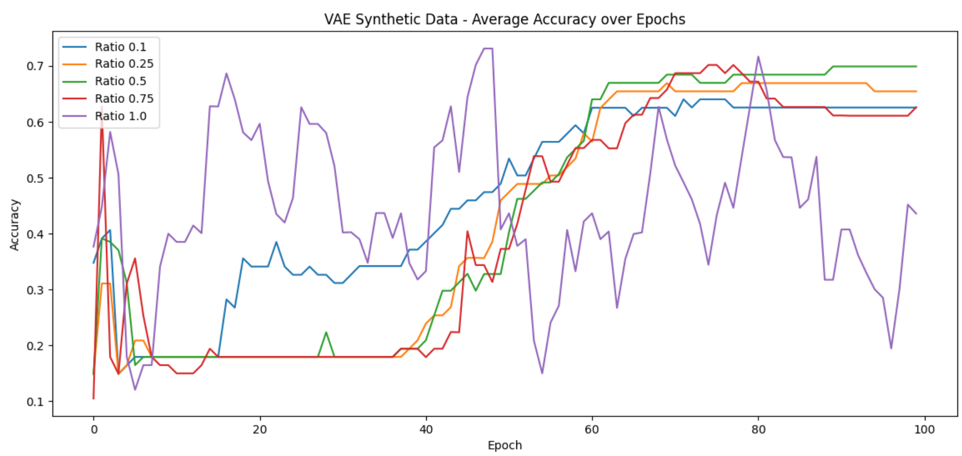

(b) Average accuracy vs. epochs plots of the classification model for different ratios of synthetic data produced using VAE training and real data.

Figure 2: Comparative performance analysis of the classification model with varying ratios of GAN and VAE-generated synthetic data and real data.

4.2 Result Analysis and Key Takeaways

The evaluation of the Synthetic Data Tuner framework revealed valuable insights into the optimal ratio of synthetic to real data for classifying components in PCB X-ray images. The GAN-generated synthetic data performed best at a ratio of 0.75, achieving an accuracy of 82.04%. This suggests that integrating a significant amount of synthetic data can boost the model's performance by offering varied and representative samples. Likewise, the VAE-generated synthetic data demonstrated optimal performance at a ratio of 0.5, with an accuracy of 69.92%. This indicates that a well-balanced combination of synthetic and real data is advantageous for the VAE-created dataset, possibly due to the superior quality of synthetic samples generated by the VAE. Various evaluation metrics such as precision, recall, and F1 score support this discovery, displaying enhancements throughout different epochs (as depicted in Figure 2). The use of stratified k-fold cross-validation during training ensured strong and dependable outcomes, emphasizing the effectiveness of synthetic data in enriching real datasets. These results highlight the potential of leveraging synthetic data to enhance model accuracy, resilience, and generalization in component classification tasks.

5. DISCUSSIONS

5.1 Strengths and Limitations

The Synthetic Data Tuner framework demonstrates various strengths, particularly its capability to utilize synthetic data for significant enhancement of deep learning models in component classification on PCB X-ray images. Through the use of advanced generative models such as GANs and VAEs, the framework produces top-notch synthetic data that enhances the training dataset, thereby improving model accuracy and robustness. The incorporation of stratified k-fold cross-validation guarantees trustworthy and impartial evaluation outcomes. Nonetheless, there are constraints to this method. Relying on the generation of high-quality synthetic data means that the initial training of GANs and VAEs must be handled with great care, as low-quality synthetic data could negatively impact the overall model performance. Furthermore, the substantial computational resources needed for training these generative models may not be feasible for all research settings.

5.2 Challenges

Several challenges were encountered during the experiments. One principal challenge involved the assurance of quality and diversity in the synthetic data produced by the GANs and VAEs. It was essential to maintain a balanced ratio between synthetic and real data, as incorrect ratios could result in overfitting or underfitting. Additionally, the computational demands of training the models posed another challenge, necessitating substantial GPU resources and time, especially for optimizing hyperparameters and carrying out cross-validation. Furthermore, the management of variations in synthetic data quality across different epochs required diligent monitoring and adaptive training strategies to uphold model stability and performance.

5.3 Future Research Directions

Future research could be directed towards several key areas in order to enhance the Synthetic Data Tuner framework and its applications. It will be essential to broaden the dataset by incorporating a wider range of PCB designs and more intricate scenarios to enhance the generalization abilities of the models. Furthermore, the integration of sophisticated techniques like semi-supervised learning and active learning might serve to further diminish the reliance on substantial amounts of annotated data. The exploration of applying the framework to alternative domains, such as medical imaging or automated inspection in manufacturing, has the potential to offer valuable insights and expand its influence. The ongoing progress in computational resources and generative model algorithms will continue to enhance the quality and efficacy of synthetic data, thus advancing the capabilities of deep learning in industrial settings.

6. CONCLUSION

This research presents the Synthetic Data Tuner framework, which effectively enhances deep learning model performance for PCB X-ray inspection by leveraging synthetic data generated from GANs and VAEs. The findings demonstrate that optimal ratios of synthetic to real data significantly improve classification accuracy, with the best ratios being 0.75 for GANs and 0.5 for VAEs. These results underscore the critical role of synthetic data in augmenting training datasets, thereby enhancing model robustness and generalization. The study highlights the importance of continued research in PCB X-ray inspection to refine these techniques further and adapt them to more diverse and complex scenarios. Practically, this framework offers substantial implications for the industry by reducing the reliance on extensive manual annotations and enabling more efficient inspection processes. Moreover, the adaptability of this approach opens opportunities for extending the framework to other domains, such as medical imaging and automated industrial inspections, thereby broadening its impact and utility.

REFERENCES

[1] Asadizanjani, N., Tehranipoor, M., and Forte, D., "Pcb reverse engineering using nondestructive x-ray tomography and advanced image processing," *IEEE Transactions on Components, Packaging and Manufacturing Technology* **7**(2), 292–299 (2017).

[2] Asadizanjani, N., Rahman, M. T., and Tehranipoor, M. M., "Physical inspection of printed circuit boards," (2021).

[3] Ghosh, S., Roy, A., Al Hasan, M. M., Craig, P., Varshney, N., Koppal, S. J., and Asadizanjani, N., "Demystifying edge cases in advanced ic packaging inspection through novel explainable ai metrics," in [*2024 IEEE 74th Electronic Components and Technology Conference (ECTC)*], 2286–2293 (2024).

[4] Khan, M. S. M., Xi, C., Varshney, N., Khan, A. A., Serna, A., Dalir, H., Sorger, V., and Asadizanjani, N., "$cm_{\mathrm{x-ray}}$: An x-ray compatibility metric for advanced packages to facilitate design-for-inspection," in [*2023 IEEE Physical Assurance and Inspection of Electronics (PAINE)*], 1–7 (2023).

[5] Mehta, D., True, J., Dizon-Paradis, O. P., Jessurun, N., Woodard, D. L., Asadizanjani, N., and Tehranipoor, M., "FICS PCB x-ray: A dataset for automated printed circuit board inter-layers inspection." Cryptology ePrint Archive, Paper 2022/924 (2022). https://eprint.iacr.org/2022/924.

[6] Roberts, A., True, J., Jessurun, N., and Asadizanjani, N., "An overview of 3d x-ray reconstruction algorithms for pcb inspection," (2020).

[7] Hasan, M. M. A., Jessurun, N., Varshney, N., and Asadizanjani, N., "Exploring the effect of annotation quality on pcb component segmentation," in [*International Symposium for Testing and Failure Analysis*], **84741**, 136–144, ASM International.

[8] Shorten, C. and Khoshgoftaar, T. M., "A survey on image data augmentation for deep learning," *Journal of Big Data* **6**, 1–48 (2019).

[9] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. C., and Bengio, Y., "Generative adversarial networks," *Communications of the ACM* **63**, 139 – 144 (2014).

[10] Kingma, D. P. and Welling, M., "Auto-encoding variational bayes," *CoRR* **abs/1312.6114** (2013).

[11] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., "Pytorch: An imperative style, high-performance deep learning library," *Advances in Neural Information Processing Systems* **32**, 8024–8035 (2019).

[12] Marcel, S. and Rodriguez, C., "Torchvision the machine-vision package," (2010).

[13] Python Imaging Library (PIL) Contributors, "Python imaging library (pil)," (1995–2011). Version 1.1.7.

[14] Radford, A., Metz, L., and Chintala, S., "Unsupervised representation learning with deep convolutional generative adversarial networks," *CoRR* **abs/1511.06434** (2015).

[15] Arjovsky, M., Chintala, S., and Bottou, L., "Wasserstein gan," (2017).

[16] Li, J., Niu, K., Liao, L., Wang, L., Liu, J., Lei, Y., and Zhang, M., "A generative steganography method based on wgan-gp," in [*Artificial Intelligence and Security*], Sun, X., Wang, J., and Bertino, E., eds., 386–397, Springer Singapore, Singapore (2020).

[17] Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y., "Spectral normalization for generative adversarial networks," (2018).

[18] He, K., Zhang, X., Ren, S., and Sun, J., "Deep residual learning for image recognition," (2015).

[19] Cui, S. and Jiang, Y., "Effective lipschitz constraint enforcement for wasserstein gan training," in [*2017 2nd IEEE International Conference on Computational Intelligence and Applications (ICCIA)*], 74–78 (2017).

[20] Ji, S., Zhang, Z., Ying, S., Wang, L., Zhao, X., and Gao, Y., "Kullback–leibler divergence metric learning," *IEEE Transactions on Cybernetics* **52**(4), 2047–2058 (2022).