https://doi.org/10.2352/EI.2024.36.6.IRIACV-277

Unsupervised PCB Anomaly Segmentation with Foundational Models

***Chih-Hui Ho*[1], *SungBal Seo*[2] *NaYeon Kim*[2], *Pin-Ying Wu*[1] *YouSuk Bae*[2], *Nuno Vasconcelos*[1]**
[1] *Department of Electrical and Computer Engineering, University of California San Diego, U.S.*
[2] *Department of Computer Engineering, Tech University of Korea, Korea*

Abstract
PCB defect segmentation aims to localize the defects in printed circuit boards (PCBs). While this problem has a great industrial impact, few datasets are publicly available. It is also challenging to predict the defects that appear during manufacturing. To address the former challenge, we curate a large dataset with various defective categories of imbalanced distribution, reflecting real-world conditions. The problem of unsupervised PCB anomaly segmentation (UAS), where no labeled defect data is available during training, is then investigated. We propose an efficient prompt tuning method to address PCB-UAS. Specifically, a pretrained large foundational segmentation model (SAM) is adapted to PCB-UAS by the introduction of a few learnable adapter layers. SAM is frozen during training and only the additional adapter parameters are learned. To overcome the lack of labeled defect images for training, we propose to create synthetic defect images that mimic the real ones. Experiments highlight that the proposed method can outperform baselines by 7 points with 16.6 times less learnable parameters.

Keyword: PCB Defect, Anomaly Segmentation, Large Foundational Models, Industrial Defect Segmentation

Introduction
Detecting and segmenting defects in industrial settings is a crucial challenge for advancing automation in various industries, often applied to electronics and PCB manufacturing [1, 4, 7, 2, 20, 23, 16]. By detecting and segmenting the PCB defects, we can identify defects in circuitry, soldering defects, misalignments, or any irregularities in the assembly process, which includes inspection for shorts, opens, or any other anomalies that might affect the functionality or reliability of the electronic components.

While this is a task of wide interest in the industry, very few labeled datasets are available due to the large variation in PCBs and their defects, making labeling defect locations costly. In addition, prior works [29, 17] also assume that defect categories are balanced distributed, which is impractical for real-world scenarios. To address these problems, we curated a novel dataset of PCB images with diverse defects and imbalanced distributions. Furthermore, to better align with the industrial setting, where segmentation labels are challenging to obtain, the proposed dataset contains only the labeled test set for evaluation, but not the training set.

To address the lack of labeled data during training, we investigate the problem of PCB unsupervised anomaly segmentation (UAS) in this work. We first synthesize images with pseudo-defect using prior knowledge of PCB defects. This is implemented by pasting these synthesized defects, including lines, dots,

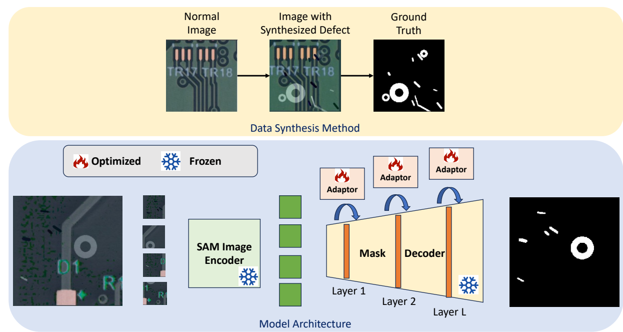

Figure 1. *Overview of the proposed PROS approach. The top part of the figure shows the data synthetic pipeline, and the bottom part shows the model architecture of the proposed PROS network.*

and circles, on normal PCB images, as shown in Figure 2. Since these pseudo-defects are synthesized, the ground truth segmentation map is automatically available, and no further annotation is needed. Note that the pseudo-defects are generated only during the training stage.

With the availability of pseudo-defects, the UAS task then becomes a supervised task, which could be solved with standard segmentation models [26, 11]. However, we observed that it does not address the UAS task by training the standard segmentation models from scratch or pre-training the segmentation models on existing datasets (like ImageNet [14]). As a result, we explore the possibility of using large foundation models like Segment Anything (SAM) [21] for this PCB defect segmentation task. SAM [21] is a foundation model for image segmentation, which is pretrained on millions of masks and can segment all objects in a scene or segment an object given a query. Since SAM is pretrained on millions of mask, we hypothesize that it can provide generalizable representations and can be adapted to UAS using few extra learnable module. To examine this hypothesis, we propose a novel and efficient architecture to segment pseudo-anomalies using the generalized feature from SAM. This is implemented by inserting a few learnable adapters between SAM's decoder layers, as shown in Figure 1. More specifically, we only optimize the adapters while keeping other parts of the SAM model frozen. We refer to the proposed framework as Pcb pROmpt Sam (PROS). Since PROS only contains few parameters, it can be optimized on the synthesized pseudo-anomalies. Experiment shows that PROS can generalize to real-world defects, which are never observed during training.

In conclusion, this work has three contributions as follows.

First, we explored the topic of PCB UAS, which is crucial for industrial application, but received limited attention in the existing literature. Second, we proposed a novel dataset featuring a broader range of defects and a more realistic distribution, mirroring real-world scenarios better than earlier datasets. Third, we proposed an innovative and effective framework PROS aimed at refining the SAM foundational segmentation model using synthesized defects. Finally, experiments demonstrate the efficiency and effectiveness of our proposed method on real defects.

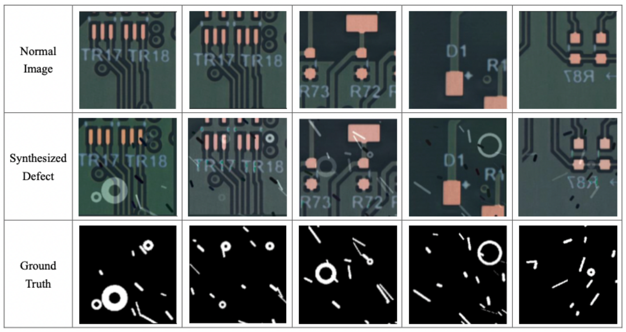

Figure 2. Illustration of the normal image, synthesized defects, and the corresponding ground truth.

Related Work

In this section, we discuss prior works on defect datasets and unsupervised anomaly segmentation methods.

Defect Datasets

Defect segmentation has been a crucial task in the industry to ensure the quality and reliability of manufacturing. Accurate segmentation allows for targeted analysis and correction of specific defect types, improving production efficiency. Multiple public defect datasets [5, 6, 39, 31, 19, 3, 37] are proposed in the literature over the past few years and covers different categories, including daily objects [5, 6, 39] and texture [31]. Unlike other defect categories, there are less PCB defect segmentation dataset [1, 4, 7, 2, 20, 23, 16]. Moreover, most of the prior datasets contains balanced sample across classes, which is not commonly seen in the real world applications. As shown in Table 1, we are the first work to consider the skew data distribution for PCB defect dataset.

Unsupervised Anomaly Segmentation (UAS)

The goal of UAS is to localize the defect in the image during inference, without seeing any defective image during training. Early works [27, 12, 10, 22, 38, 13, 30, 18] can be coarsely categorized into three categories, including student-teacher methods, flow based methods and reconstruction based methods. **Student-teacher methods** [38, 13, 30] contain a pretrained teacher encoder (i.e. pretrained on ImageNet) and a student encoder. During training, both the teacher encoder and student encoder consumes the normal images and the difference between their output are minimized. During inference, the prediction differences between the student and the teacher encoder suggests the anomaly region. **Flow based methods** [33, 28, 25] aim to regularize the distribution of encoder output, when consuming the normal training

	# of Images	# of Categories	Color	Imbalanced
DeepPCB [29]	8853	6	No	
Peking-PCB [17]	1386	6	Yes	
Ours (PROS)	10261	10	Yes	V

Table 1. Comparison of the prior datasets with the proposed dataset.

images. More specifically, the distribution of encoder output are encouraged to be, for example, a Gaussian distribution using normalizing flow [25]. During inference, the encoder output of an abnormal image will be out-of-distribution. **Reconstruction based methods** [35, 24, 36, 34, 32] seeks to project the input image (both normal and abnormal image) to the normal image manifold. This is implemented by minimizing the reconstruction difference between input normal image and the reconstruction output during training. During inference, the region with large reconstruction differences are considered as anomalies. Unlike these earlier works, we adopt the large segmentation model for the UAS task.

Method

In this section, we propose the PROS framework for the PCB UAS task. PROS is a lightweight framework optimized on the pseudo-anomalies. The entire framework is illustrated in Figure 1, and more details are provided below.

Architecture

To address the UAS problem, we propose to leverage the generalized feature from a large foundational model SAM [21]. SAM contains an image encoder, a prompt encoder, and a mask decoder. The image encoder extracts the image embedding using the vision transformer [15] . The mask decoder is based on DETR [8] and MaskFormer [9], to predict semantic and instance-level segmentations. Please refer to the SAM [21] paper for additional details.

To leverage the powerful pre-trained knowledge in SAM, we insert K learnable adaptation module to SAM. More specifically, each learnable adaptation module is a fully connected layer and is inserted at each transformer layer of SAM decoder. During training, we freeze the majority part of the SAM model and only tune the additional parameters. This efficient design allows SAM to detect PCB defects with only a few learnable parameters.

Training with Synthetic Data

Gathering a large dataset with annotated information is impractical in industrial settings for two key reasons. Firstly, defects in the manufacturing industry are inherently rare occurrences, making it challenging to assemble a sufficiently large and diverse labeled dataset. Secondly, the annotation process requires domain expertise, often unavailable on public crowd-sourcing platforms commonly used for labeling extensive datasets. These challenges emphasize the difficulty in obtaining labeled data at scale for industries dealing with infrequent defects and specialized knowledge requirements.

To overcome these challenges, we employed the prior knowledge that the majority of defects on PCBs are dots, lines, and circles. We synthesized these defects by using OpenCV and pastes these defects onto regular images. During the pasting, we blur these synthesized defect with different transparency. Since these pseudo defects are synthesized, this allows us to derive the positions and categories of the defects automatically. As shown in

Architecture	Pretrained Dataset	Training Loss	# of Learnable Parameters (M)	Mean IOU
UNet	N/A	BCE	31	38.24
UNet	MRI Images	BCE	7.8	0.85
UNet	ImageNet	BCE	14.3	39.52
Siamese UNet	N/A	BCE	42.8	38.89
Siamese UNet Diff	N/A	BCE	31	39.08
UNet	N/A	Soft Dice	31	38.26
UNet	MRI Images	Soft Dice	7.8	0.87
UNet	ImageNet	Soft Dice	14.3	38.17
Siamese UNet	N/A	Soft Dice	42.8	39.05
Siamese UNet Diff	N/A	Soft Dice	31	39.13
UNet	N/A	Cross Entropy	31	39.07
UNet	MRI Images	Cross Entropy	7.8	10.3
UNet	ImageNet	Cross Entropy	14.3	38.38
Siamese UNet	N/A	Cross Entropy	42.8	38.66
Siamese UNet Diff	N/A	Cross Entropy	31	39.15
Ours (PROS)	SAM Dataset	Cross Entropy	0.47	46.02

Table 2. Comparison of the proposed method PROS with prior works. PROS requires fewer parameters to learn and achieves better performance.

Figure 2, the ground truth is available without additional manual annotation.

Experiment

This section discusses the details of the dataset and presents the experimental results of the proposed PROS on the UAS task.

Dataset and Metric

Our proposed dataset contains 10261 images, which comprises 10 defect types with imbalanced distribution as shown in Figure 4. To analyze the dataset distribution, we adopt the imbalanced factor β, which is defined as $\beta = \frac{max_c\{N_c\}}{min_c\{N_c\}}$. For our dataset, the imbalance factor is around 271, which is skewer than most of the imbalanced dataset considered in the long-tailed literature.

To measure the performance of PROS and the baselines, we adopt the dice score, which is defined as

$$DICE = \frac{2|GT \cap PRED|}{|GT| + |PRED|}, \qquad (1)$$

where GT is the ground truth defect segmentation and $PRED$ is the predicted defect area.

Implementation Details

To train PROS, we use the pretrained SAM model from its official github. More specifically, we use the SAM with the Vit-Base image encoder. Each input image is scale to 256×256 before feed into the SAM encoder. We trained PROS with Pytorch for 10 epochs using batch size of 4. We use AdamW optimizer and the learning rate is set to 0.0001. The entire training is conducted on the a Titan Nvidia GPU.

Quantitative Results

We first introduce a set of baselines for the PCB-UAS task. Both the baselines and PROS are trained only on synthesized defect images. As shown in Figure 3, we consider the UNet [26], Siamese UNet, and Siamese UNet Diff [11] architectures. For

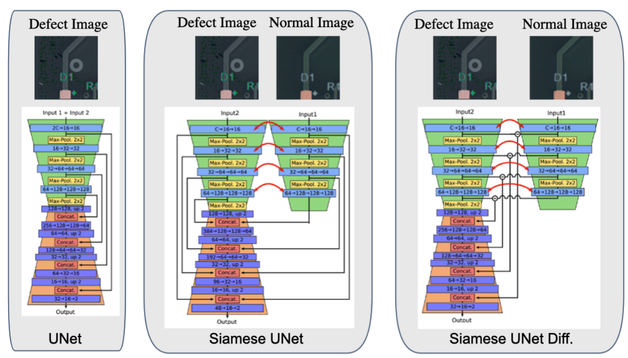

Figure 3. *Baseline architecture. (Left) UNet can be trained from scratch or initialized with weight pretrained on ImageNet images or MRI images. (Middle) Siamese UNet takes both the normal image and the defect image as input. The image features of the normal image and defect image are concatenated and passed to the UNet decoder. (Right) Siamese UNet Diff. is similar to Siamese UNet, but it uses the difference of intermediate features of two images for segmentation.*

the UNet baselines, it can be trained from scratch or pretrained on MRI or ImageNet datasets. Three different losses are used for each baseline, including the binary cross entropy (BCE) loss, soft dice loss and the cross entropy loss.

Table 2 summarizes the benefits of the proposed PROS approach. The table compares the dice score between the baselines and PROS on the labeled test set of the proposed dataset. PROS significantly outperforms the baselines by about 7 points. Furthermore, since PROS is built upon the foundational model SAM, it requires only 0.47M learnable parameters, which is 16.6 times less than the baselines.

Ablations on Number of Adapters

We conduct ablation experiments by varying the number of layers for inserting the learnable adapters. As shown in Table 3,

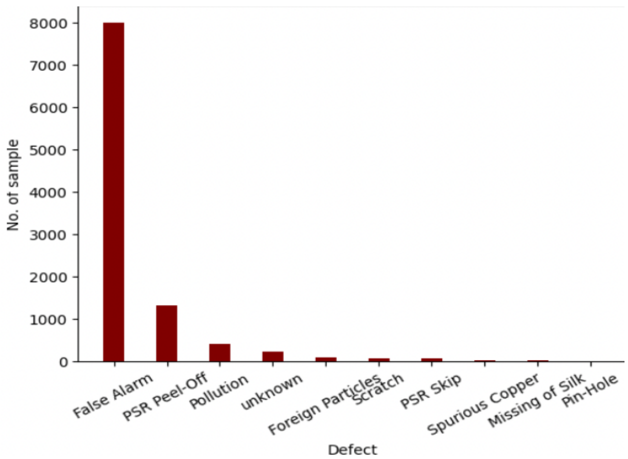

Figure 4. *Imbalance sample distribution of the proposed dataset across different defect categories.*

	Insert K adapters					
	K=12	K=11	K=9	K=7	K=4	K=1
# of Param. (M)	3.76	3.46	2.86	2.26	1.37	0.47
Mean IOU	21.78	30.46	34.36	35.87	42.7	46.02

Table 3. Ablation of the mean IOU and number of learnable parameters with respect to the number of adapters.

the weakest performance is observed when adapters are inserted into all decoder layers of SAM (i.e. $K = 12$). By reducing the number of adapters, the performance of PROS increases. This suggests that the model tends to overfit on the synthesized defects for large K and cannot generalize to real defects.

Qualitative Results

Figure 5 shows some quantitative results of the proposed PROS on these four input images. Note that these four images are sampled from the test set, which contains real defects. The results show the strong generalizability of PROS on real defects.

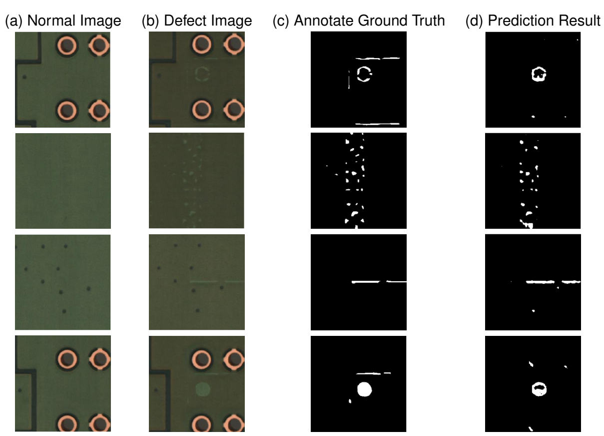

Figure 5. *Qualitative results of PROS on the real defects, which are not seen during training.*

Conclusion and Future Work

In this work, we consider the problem of PCB anomaly segmentation with unsupervised scenarios (PCB-UAS), which has a great industrial impact. We first curate a large dataset with various defective categories of imbalanced distribution, reflecting real-world conditions. The problem of unsupervised PCB anomaly segmentation (UAS) is investigated, where no labeled defect data is available during training. We propose an efficient fine-tuning method to address PCB-UAS. Specifically, a pre-trained large foundational model (SAM) is adapted to PCB-UAS by introducing a few learnable adapter layers. SAM is frozen during training, and only the additional adapter parameters are learned. To overcome the lack of labeled defect images for training, we propose to create synthetic defect images that mimic the real ones. Experiments highlight that the proposed method can outperform baselines by 7 points with 16.6 times less learnable parameters. Comprehensive experiments showcase the effectiveness of the proposed framework. We expect this work could pave the way for a novel research direction in PCB defect detection with limited annotation.

Acknowledgement

This work was supported by Innovative Human Resource Development for Local Intellectualization program through the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (IITP-2024-2020-0-01741).

References

[1] Venkat Anil Adibhatla, Huan-Chuang Chih, Chi-Chang Hsu, Joseph Cheng, Maysam F Abbod, and Jiann-Shing Shieh. Defect detection in printed circuit boards using you-only-look-once convolutional neural networks. *Electronics*, 9(9):1547, 2020.

[2] Yolanda D Austria and Arnel C Fajardo. Defect detection and classification in printed circuit boards using convolutional neural networks. In *2023 2nd International Conference on Edge Computing and Applications (ICECAA)*, pages 1498–1504. IEEE, 2023.

[3] Haoping Bai, Shancong Mou, Tatiana Likhomanenko, Ramazan Gokberk Cinbis, Oncel Tuzel, Ping-Chia Huang, Jiulong Shan, Jianjun Shi, and Mengsi Cao. Vision datasets: A benchmark for vision-based industrial inspection. *ArXiv*, abs/2306.07890, 2023.

[4] Xiaolong Bai, Yuming Fang, Weisi Lin, Lipo Wang, and Bing-Feng Ju. Saliency-based defect detection in industrial images by using phase spectrum. *IEEE Transactions on Industrial Informatics*, 10(4):2135–2145, 2014.

[5] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTec AD — A comprehensive real-world dataset for unsupervised anomaly detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9584–9592, 2019.

[6] Paul Bergmann, Xin Jin, David Sattlegger, and Carsten Steger. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *ArXiv*, abs/2112.09045, 2021.

[7] YR Bhanumathy, MP James, Shivangi Jha, and Sudeesh Balan. Defect detection in pcbs using convolutional neural network. In *2021 International Conference on Recent Trends on Electronics, Information, Communication & Technology (RTEICT)*, pages 382–386. IEEE, 2021.

[8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.

[9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel

classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021.

[10] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *ArXiv*, abs/2005.02357, 2020.

[11] Rodrigo Caye Daudt, Bertr Le Saux, and Alexandre Boulch. Fully convolutional siamese networks for change detection. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 4063–4067. IEEE, 2018.

[12] Thomas Defard, Aleksandr Setkov, Angélique Loesch, and Romaric Audigier. PaDiM: A patch distribution modeling framework for anomaly detection and localization. In *ICPR Workshops*, 2020.

[13] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9727–9736, 2022.

[14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[16] Rudi Heriansyah, Syed Abdul Rahman Syed Abu Bakar, and Muhammad Mun'im Ahmad Zabidi. *Segmentation of PCB Image Into Simple Generic Patterns Using Mathematical Morphology and Windowing Technique*. Universiti Teknologi Malaysia, 2002.

[17] Weibo Huang and Peng Wei. A pcb dataset for defects detection and classification. *arXiv preprint arXiv:1901.08204*, 2019.

[18] Jeeho Hyun, Sangyun Kim, Giyoung Jeon, Seungwook Kim, Kyunghoon Bae, and Byungjin Kang. ReConPatch: Contrastive patch representation learning for industrial anomaly detection. *ArXiv*, abs/2305.16713, 2023.

[19] Stepan Jezek, Martin Jonak, Radim Burget, Pavel Dvorak, and Milos Skotak. Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions. In *2021 13th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, pages 66–71, 2021.

[20] Jungsuk Kim, Jungbeom Ko, Hojong Choi, and Hyunchul Kim. Printed circuit board defect detection using deep learning via a skip-connected convolutional autoencoder. *Sensors*, 21(15):4968, 2021.

[21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.

[22] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. CutPaste: Self-supervised learning for anomaly detection and localization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9659–9669, 2021.

[23] Zhigang Ling, Aoran Zhang, Dexin Ma, Yuxin Shi, and He Wen. Deep siamese semantic segmentation network for pcb welding defect detection. *IEEE Transactions on Instrumentation and Measurement*, 71:1–11, 2022.

[24] Shancong Mou, Xiaoyi Gu, Meng Cao, Haoping Bai, Ping Huang, Jiulong Shan, and Jianjun Shi. RGI: Robust GAN-inversion for mask-free image inpainting and unsupervised pixel-wise anomaly detection. In *ICLR*, 2023.

[25] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *ArXiv*, abs/1505.05770, 2015.

[26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.

[27] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Scholkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14298–14308, 2021.

[28] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but differnet: Semi-supervised defect detection with normalizing flows. In *Winter Conference on Applications of Computer Vision (WACV)*, January 2021.

[29] Sanli Tang, Fan He, Xiaolin Huang, and Jie Yang. Online pcb defect detector on a new pcb defect dataset. *arXiv preprint arXiv:1902.06197*, 2019.

[30] Tran Dinh Tien, Anh Tuan Nguyen, Nguyen Hoang Tran, Ta Duc Huy, Soan T.M. Duong, Chanh D. Tr. Nguyen, and Steven Q. H. Truong. Revisiting reverse distillation for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24511–24520, June 2023.

[31] Matthias Wieler, Tobias Hahn, and Fred. A. Hamprecht. Weakly supervised learning for industrial optical inspection, 2007. `https://hci.iwr.uni-heidelberg.de/node/3616`.

[32] Zhiyuan You, Lei Cui, Yujun Shen, Kai Yang, Xin Lu, Yu Zheng, and Xinyi Le. A unified model for multi-class anomaly detection. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 4571–4584. Curran Associates, Inc., 2022.

[33] Jiawei Yu1, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. FastFlow: Unsupervised anomaly detection and localization via 2D normalizing flows. *ArXiv*, abs/2111.07677, 2021.

[34] Vitjan Zavrtanik, Matej Kristan, and Danijel Skocaj. DRAEM - A discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8330–8339, October 2021.

[35] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-GAN: High-fidelity defect synthesis for automated defect inspection. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2523–2533, 2021.

[36] Hui Min Zhang, Z. Wang, Zuxuan Wu, and Yuwei Jiang. DiffusionAD: Denoising diffusion for anomaly detection. *ArXiv*, abs/2303.08730, 2023.

[37] Jian Zhang, Runwei Ding, Miaoju Ban, and Ge Yang. PKU-GoodsAD: A supermarket goods dataset for unsupervised anomaly detection and segmentation. *ArXiv*, abs/2307.04956, 2023.

[38] Xuan Zhang, Shiyu Li, Xi Li, Ping Huang, Jiulong Shan, and Ting Chen. DeSTSeg: Segmentation guided denoising student-teacher for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3914–3923, June 2023.

[39] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. *arXiv preprint arXiv:2207.14315*, 2022.