Polar-Ano: Surface Anomaly Detection via Deep Polarization Imaging and Data Synthesis with Physic-based Rendering

Yaxu Xie, Alain Pagani, and Didier Stricker
Augmented Vision
German Research Center for Artificial Intelligence
Kaiserslautern, Germany
name.surname@dfki.de

***Abstract*—In this paper, we propose a complete pipeline from generating polarization images via physics-based rendering to train and deploy an image anomaly detection and localization model for polarimetric industrial inspection. The method consists of two stages. We first compute the Polarimetric Priors with both determined and learning-based method. Then, the Polarimetric Priors are given to a self-supervised surface anomaly detection network to predict the anomalies score and anomalies masks. To train the network, we adapt and modify a physic-based rendering pipeline to generate photo-realistic data samples of polarized images on a large scale. Our experiments show the effectiveness of our proposed pipeline.**

***Index Terms*—anomaly detection, polarimetric imaging, physic-based rendering, shape-from-polarization**

I. Introduction

Surface anomalies frequently arise due to changes in material composition, surface damages, and stains, or structural damages and deformities. The detection of surface anomaly requires not only binary classification of normal and anomaly samples but also the localization of the anomaly in image space. In the domain of industrial anomaly detection, the pursuit of enhancing sensitivity and accuracy has driven the exploration of novel techniques for image appearance acquisition that can discriminate various and subtle anomalies from expected patterns. However, many types of industrial products are usually texture-less, single-colored, or even transparent, which may elude conventional imaging techniques (simple RGB or gray-scaled images).

Polarimetric imaging is a well-known technology that capitalizes on the polarization properties of light and provides distinctive information about object surfaces. By extending the analysis beyond calssical intensity-based imaging, polarimetric imaging has demonstrated the potential to unearth latent anomalies that might otherwise remain concealed. The surface geometric, roughness, texture, and micro-structures can be revealed, via analyzing the polarimetric information such as degree of linear polarization (DoLP), angle of linear polarization (AoLP), and stoke maps, which contributes to

The research leading to these results has been partially funded by the German BMBF project RACKET (01IW20009).

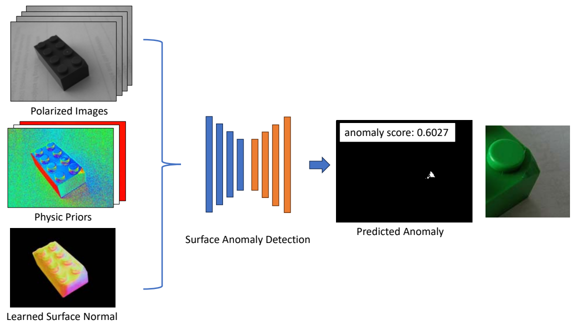

Fig. 1: **Surface anomaly detection with multiple inputs**, including polarized images, physic priors, and learned surface normal.

the identification of anomalies. Although the reflection and polarization of light on specific material are well-studied problems and can be solved using Fresnel equation, there are still several limitations of traditional determined polarimetric imaging: (1) The acquisition of polarimetric data is inherently prone to Poisson noise since the intensity of light through linear polarizer is reduce by 50%, (2) Polarimetric imaging often confronts the ambiguity inherent in interpreting reflections from complex surfaces, (3) The surface of real-world objects can consist of multiple materials or have different roughness, which again brings ambiguity to the model [30]. Since deep learning methods show a strong ability to learn priors and to be robust against noise in many general domains, researchers have noticed its potential in polarimetric imaging [4], [17].

Yet another challenge is that anomaly appearances are very rare compared to the normal samples, which can be laborious and costly to obtain in real-world industrial settings. Supervised learning methods require extensive labeled datasets for training, but the imbalanced data distribution limited its performance. In response, self-supervised deep learning emerges as a promising pipeline to address these challenges, offering a data-efficient and resource-conscious approach to

979-8-3503-7091-1/23/$31.00 ©2023 IEEE

DOI 10.1109/SITIS61268.2023.00055

surface anomaly detection. By reducing the dependency on annotated datasets and enhancing the model's adaptability, self-supervised learning contributes to more cost-effective and robust anomaly detection solutions. Although the self-supervised anomaly detection method requires only normal data samples for training, the intensive labor of manually creating and labeling polarimetric data samples is still a bottleneck for combining polarimetric imaging and learning-based surface anomaly detection.

Therefore, we adopt a physics-based rendering pipeline to automatically generate high-quality polarized image samples of target 3D objects with the desired ground truth. This pipeline allows us to introduce controlled randomness of light intensity and direction, object materials, color and texture, and background to the synthetic data samples. Then, we develop a vision system for surface anomaly detection using polarimetric imaging and self-supervised learning on our generated data. The contributions of our paper are the following:

- we implemented a physics-based rendering pipeline to generate polarization images of object-level scene layouts, which enables us to train and test learning-based methods on it.
- we proposed the comprehensive use of multiple inputs (raw image, DoP, AoP and surface normal) to enrich the appearance information.
- we trained a self-supervised surface normal detection method on multiple inputs, and test the performance of the model on both synthetic and real data samples.

II. Related Work

In this section, we first review the recent progress of polarimetric imaging, especially learning-based ones. Then we give a summary of industrial surface anomaly detection. Finally, we discuss the existing pipelines of physics-based rendering for polarization.

Polarimetric Imaging has been strongly boosted recently with the progress of on-chip polarization CMOS, which enables to capture of images with multiple different polarization angles in one shot. Polarization cues of objects and scenes, such as Degree of Polarization (DoP), Angle of Polarization (AoP) and stroke vectors can be computed with the polarized images and provide extra information for various tasks, such as surface normal estimation [14], [30], spatially varying surface reflectance functions (SVBRDF) recovery [5], reflection separation [22] and anomaly detection [6]. However, the ambiguities of polarization cues due to unknown reflection types, object surface materials, and sensor noises limit the application scenarios of many works. Recently, deep learning has been introduced into polarimetric imaging for the reason that neural networks are more adaptive facing ambiguities and imperfection of input data. Learning-based Shape-from-Polarization (SfP) has been tackled in both object [4] and scene [18] levels as well as the acquisition of SVBRDF [10]. Sparse polarization sensor and the polarization information compensation network [17] are also developed, to overcome the sensor sensitivity reduction resulting from polarization filters. The pose estimation and anomaly detection of transparent objects are also explored in several works [12], [16], [32].

Self-supervised Surface Anomaly Detection In industrial quality control, the surface anomaly appearances are significantly diverse. Moreover, the images of anomaly samples are too rare to be manually collected/created and annotated in time-cost efficient ways. Self-supervised learning methods are therefore widely studied in anomaly detection. To enable self-supervised training, some works [2], [3], [7] utilize auto-encoders and generative adversarial networks to reconstruct normal data. These methods are fully trained with normal data, during inference the anomalies are identified according to the reconstruction quality, which means the network is supposed to show lower reconstruction performance in anomaly regions. This assumption may fail if the reconstruction network is too representative or the anomalies are too similar to normal appearance for photometric loss. Methods like $DRÆM$ [33] and CutPaste [19] make efforts to generate and blend random synthetic anomalies on anomaly-free images and using a discriminator network to compare the input and the reconstructed image so that the anomaly can be detected. Other works [8], [25], [26] have noticed that the normal features and the anomalous features can be normalized and clustered far away from each other in the embedding space. SimpleNet [21] is then proposed to generate anomaly in embedding space and perform transfer learning to reduce the distribution gap between pretrained and target datasets, which results in a lightweight and fast design.

Data Synthetic using Physic-based Rendering Synthetic data is generated via rendering methods which means it does not require manual data collection efforts and it can contain nearly perfect annotations. Many large-scaled synthetic datasets are already widely used in nowadays learning-based computer vision research, such as Replica [28] and Hypersim [24] for indoor spatial sensing, Virtual KITTI [11] and SIFT [29] for autonomous driving. To reduce the gap between synthetic and real data, physically based rendering (PBR) software is used during the generating of these datasets, which aims to simulate light interaction with materials in a way that accurately represents the real world. Polarization rendering is a specialized aspect of PBR that focuses on capturing light polarization effects. Early research [23], [27], [31] already established the model and the numeric solver of polarization rendering on complex/mixed reflection situations of different materials. Recently, efforts have also been made by Wenzel Jakob et al. [15] to develop a research-oriented rendering system for forward and inverse light transport simulation, which is differentiable, cross-platformed, and adaptive to modern CPU and GPU acceleration backend.

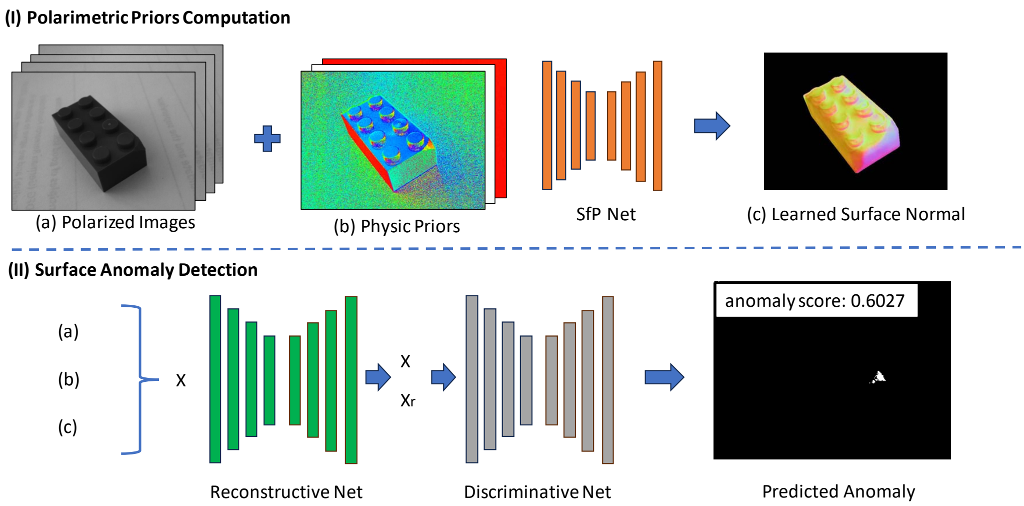

Fig. 2: **Network Architecture.** The network consists of two stages. At the first stages, the Polarimetric Priors X are computed. And then the Polarimetric Priors are given to the anomaly detection network to predict the anomalies.

III. Approach

In this section, we look into the details of our proposed polarimetric surface anomaly detection pipeline. We first give definitions of every polarimetric priors used as the input of the surface anomaly detection method. After clarifying the architecture and loss terms of the network, we explain the data synthesis method we adopted to render polarization images. The overview of the network architecture is shown in Figure 2.

A. Polarimetric Prior Estimation

The polarization state of light can be affected by the reflection. When unpolarized light strikes a surface, the reflected light can become partially polarized. And the azimuth and zenith angle of the reflection are determined by the surface normal of the observed object. The relation between polarization, reflection and object surface geometry can be jointly described by Fresnel Equation. Therefore, estimating the polarimetric information will provide extra cues for identify the surface anomaly.

We assume the the polarization images $\mathbf{I}(\phi_{pol})$ are from a polarization camera with on-chip lens and polarizers with polarization angles $\phi_{pol} \in \{0°, 45°, 90°, 135°\}$, which can be determined by:

$$\mathbf{I}(\phi_{pol}) = \mathbf{I}_{un}(1 + \rho\cos(2\phi - 2\phi_{pol})) \tag{1}$$

where ϕ is the angle of polarization (AoP), ρ is the degree of polarization (DoP) and $\mathbf{I}_{un}$ is the unpolarized intensity of light. Having the images with different polarizer angles, the degree of polarization ρ and the angle of polarization ϕ are computed by:

$$\begin{aligned} \rho &= \frac{\sqrt{((I_0 - I_{90})^2 + ((I_{45} - I_{135})^2)}}{I}, \\ \phi &= \frac{1}{2}\arctan\frac{I_{45} - I_{135}}{I_0 - I_{90}} \end{aligned} \tag{2}$$

with $I = (I_0 + I_{45} + I_{90} + I_{135})/2$.

Due to the complex and unknown material behavior and multiple solution natural of Fresnel Equation, the traditional shape-from-polarization methods are valid only under strict assumptions and perform worse than learning-based methods [4], [18]. We therefore utilize the learned surface normal from pretrained SPW-Net [18] as the third polarimetric priors. Thus, the input for our polarimetric surface anomaly detetion method consists of the unpolarized image, the degree of polarization, the angle of polarization and the learned surface normal:

$$X = \{I_{un}, \rho, \phi, n\}\,. \tag{3}$$

B. Surface Anomaly Detection

We utilize the $DRÆM$ Method [33] as our surface anomaly detection baseline. $DRÆM$ is a reconstructive-discriminative network, which can be trained in an end-to-end manner on synthetically generated just-out-of-distribution patterns. The model consist of two sub-networks, the reconstructive network takes image samples with augmented anomalies as input and reconstructs the anomaly-free images from them. While the discriminative network takes both the reconstructed images and the anomaly-augmented images as inputs and estimate the region of anomalies in image space. The synthetic anomalies are created using random picked patterns masked with a binary mask generated using Perlin noise genera- tor, and blended on the anomaly-free image. We propose to train $DRÆM$ on our Polarimetric Prior $X = \{I_{un}, \rho, \phi, n\}$ separately:

$$\hat{M}_i, \hat{\eta}_i = f(X_i) \quad with \quad X_i \in \{I_{un}, \rho, \phi, n\} \tag{4}$$

where $\hat{M}$ is the predicted anomaly mask with the same frame size as the input images, and $\hat{\eta} \in (0, 1)$ is the predicted anomaly score. We then simply compute the average of the M_i and η_i as the final output.

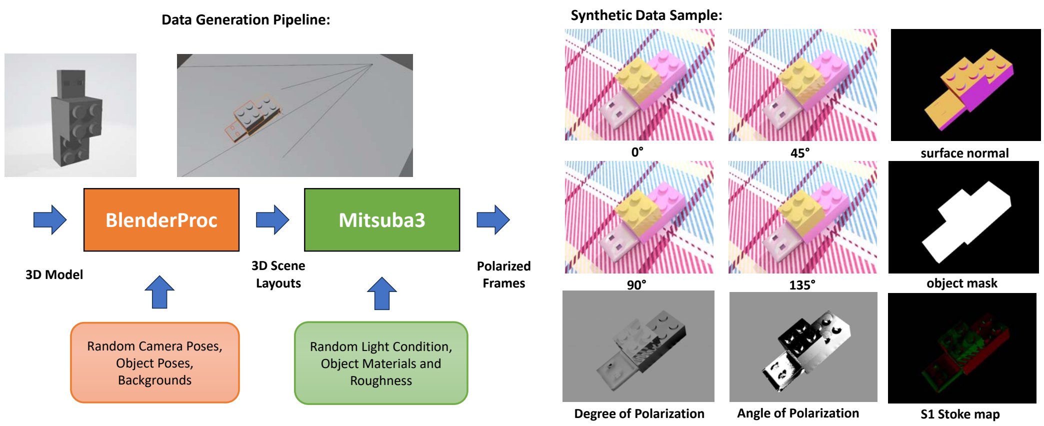

Fig. 3: **Our physics-based polarization rendering pipeline for data synthesis.** Here the reference object is a USB stick with the shape of plastic construction bricks, the top and the bottom bricks is assign with different material smoothness.

Same as the original design of [33], the reconstructive network is trained with photometric loss, while the discriminative network is trained Focal Loss [20] for the predicted anomaly mask:

$$\mathcal{L} = \mathcal{L}_{rec}(X_i, X_i^r) + \mathcal{L}_{seg}(M, \hat{M}_i) \tag{5}$$

where M is the ground truth anomaly mask, and X_i^r is the reconstructed term of X_i.

C. Physic-based Polarization Data Synthesis

Instead of capturing real polarimetric images and labeling them with ground truth object masks, surface normals, and anomaly maps, we modify an open-sourced physic-based rendering method [1] to generate synthetic data samples. As illustrated in Figure 3, our method bases on BlenderProc [9] and Mitsuba3 [15]. We first create the 3D model of target object via CAD software and generate a high resolution 3D mesh of each component of the object. Then we set up a scene layout as in the BOP dataset [13] with random sampled camera poses, backgrounds and object poses using BlenderProc.

The scene layout is later on converted to Mitsuba3 scripts to conduct physic-based rendering to generate photo-realistic polarization images. As shown in Figure 3, every component of the object (here a USB stick shaped as a plastic construction brick) is assigned with different BSDFs (polarized plastic and rough conductor).

Comparing to the original implementation [1], we also introduce controlled randomness such as spot light position, light intensity, material color etc, in order to increase the generalization-ability of our data rendered samples. All manipulable parameters in our setup is listed in Table I.

IV. Experiments

In this section, we first give details about the implementation. Then we test our trained method on our own dataset and real data samples, and evaluate the performance qualitatively and quantitatively.

Variables(Units)	Sampling Methods	Range
Camera Pose (m)	Spherical Shell	$r \in (0.4, 0.49)$
Object Pose (m)	Surface Dropping	$h \in (1, 4)$
Background	-	-
Spot Light Position (m)	Spherical Shell	$r \in (1, 1.5)$
Light Intensity (w/m^2)	Uniform	$p \in (2.25, 3.375)$
Material Color	Uniform	$i \in (0, 1)$
Material α (Rough)	Uniform	$\alpha \in (0.1, 0, 7)$
Material α (Smooth)	Uniform	$\alpha \in (0.001, 0.011)$

TABLE I: Controlled random variables of PBR data synthesis.

Implementation Details We generated 4,608 synthetic polarized data samples with annotated object mask, DoP, AoP and surface normal. Samples have a frame size of 480×640. We randomly picked 80% of the samples for training, and 20% for testing. Our network is trained on the generated data for 50 epochs, using ADAM optimizer with an initial learning rate of 0.0001 and batch size of 4. Because SPW [18] is trained on scene-level dataset, we fine-tuned the network using our synthetic dataset for 10 epochs.

Evaluation We first evaluate our method on the test split of our synthetic dataset in terms of AUROC (Area Under the ROC Convex Hull) both for image-wise anomaly (AUROC-I) and pixel-wise anomaly (AUROC-P), the results is listed in Table II. We mark the input variants with only raw unpolarized image as R, with AoP as ϕ, with DoP as ρ and with learned surface normal as n. As the result shows, by adding Polarimetric Pirors one by one, the results keep improving. With the method variant $R+\phi+\rho+n$ (all of the Polarimetric Priors) yields the best result among all variants. We also test the performance of our method on real data samples, the visualized results are given in Figure 4. Even though our method is only trained on synthetic data, it still performs well enough for real world samples, which also proves the reaslism and usefulness of our render synthetic data.

Methods	I-AUROC	P-AUROC
R	99.075	99.070
R+ϕ	99.689	99.836
R+ϕ+ρ	99.794	99.911
R+ϕ+ρ+n	**99.817**	**99.926**

TABLE II: Comparison of the performance of the surface anomaly detection method with different input variants.

V. Conclusion

A polarimetric surface anomaly detection method was presented, together with the method for generate training data via physic-based rendering techniques. Our proposed pipeline shows accurate result during testing both in synthetic and real data. For the future work, we would like to explore the rendering techniques to generate additional types of realistic anomalies to create a more comprehensive polarimetric anomaly detection benchmark.

References

[1] A-guridi. Mistubarenderer. https://https://github.com/A-guridi/MistubaRenderer, 2021.

[2] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 622–637. Springer, 2019.

[3] Samet Akçay, Amir Atapour-Abarghouei, and Toby P Breckon. Skip-ganomaly: Skip connected and adversarially trained encoder-decoder anomaly detection. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019.

[4] Yunhao Ba, Alex Gilbert, Franklin Wang, Jinfa Yang, Rui Chen, Yiqin Wang, Lei Yan, Boxin Shi, and Achuta Kadambi. Deep shape from polarization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 554–571. Springer, 2020.

[5] Seung-Hwan Baek, Daniel S Jeon, Xin Tong, and Min H Kim. Simultaneous acquisition of polarimetric svbrdf and normals. *ACM Trans. Graph.*, 37(6):268–1, 2018.

[6] Brent D Bartlett, Ariel Schlamm, Carl Salvaggio, and David W Messinger. Anomaly detection of man-made objects using spectropolarimetric imagery. In *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XVII*, volume 8048, pages 109–115. SPIE, 2011.

[7] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *arXiv preprint arXiv:1807.02011*, 2018.

[8] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pages 475–489. Springer, 2021.

[9] Maximilian Denninger, Dominik Winkelbauer, Martin Sundermeyer, Wout Boerdijk, Markus Knauer, Klaus H. Strobl, Matthias Humt, and Rudolph Triebel. Blenderproc2: A procedural pipeline for photorealistic rendering. *Journal of Open Source Software*, 8(82):4901, 2023.

[10] Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. Deep polarization imaging for 3d shape and svbrdf acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15567–15576, 2021.

[11] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4340–4349, 2016.

[12] Daoyi Gao, Yitong Li, Patrick Ruhkamp, Iuliia Skobleva, Magdalena Wysocki, HyunJun Jung, Pengyuan Wang, Arturo Guridi, and Benjamin Busam. Polarimetric pose prediction. In *European Conference on Computer Vision*, pages 735–752. Springer, 2022.

[13] Tomáš Hodaň, Frank Michel, Eric Brachmann, Wadim Kehl, Anders Glent Buch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, Caner Sahin, Fabian Manhardt, Federico Tombari, Tae-Kyun Kim, Jiří Matas, and Carsten Rother. BOP: Benchmark for 6D object pose estimation. *European Conference on Computer Vision (ECCV)*, 2018.

[14] Tomoki Ichikawa, Matthew Purri, Ryo Kawahara, Shohei Nobuhara, Kristin Dana, and Ko Nishino. Shape from sky: Polarimetric normal recovery under the sky. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14832–14841, 2021.

[15] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. https://mitsuba-renderer.org.

[16] Atsutake Kosuge, Lixing Yu, Mototsugu Hamada, Kazuki Matsuo, and Tadahiro Kuroda. A deep metric learning-based anomaly detection system for transparent objects using polarized-image fusion. *IEEE Open Journal of the Industrial Electronics Society*, 4:205–213, 2023.

[17] Teppei Kurita, Yuhi Kondo, Legong Sun, and Yusuke Moriuchi. Simultaneous acquisition of high quality rgb image and polarization information using a sparse polarization sensor. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 178–188, 2023.

[18] Chenyang Lei, Chenyang Qi, Jiaxin Xie, Na Fan, Vladlen Koltun, and Qifeng Chen. Shape from polarization for complex scenes in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12632–12641, 2022.

[19] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9664–9674, 2021.

[20] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.

[21] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. Simplenet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20402–20411, 2023.

[22] Youwei Lyu, Zhaopeng Cui, Si Li, Marc Pollefeys, and Boxin Shi. Reflection separation using a pair of unpolarized and polarized images. *Advances in neural information processing systems*, 32, 2019.

[23] Miyazaki, Tan, Hara, and Ikeuchi. Polarization-based inverse rendering from a single view. In *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 982–987. IEEE, 2003.

[24] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *International Conference on Computer Vision (ICCV) 2021*, 2021.

[25] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14318–14328, 2022.

[26] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but differnet: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1907–1916, 2021.

[27] Sairam Sankaranarayanan. *Modelling polarized light for computer graphics*. Iowa State University, 1997.

[28] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.

[29] Tao Sun, Mattia Segu, Janis Postels, Yuxuan Wang, Luc Van Gool, Bernt Schiele, Federico Tombari, and Fisher Yu. SHIFT: a synthetic driving dataset for continuous multi-task domain adaptation. In *Computer Vision and Pattern Recognition*, 2022.

[30] Vage Taamazyan, Achuta Kadambi, and Ramesh Raskar. Shape from mixed polarization. *arXiv preprint arXiv:1605.02066*, 2016.

[31] Lawrence B Wolff and David J Kurlander. Ray tracing with polarization parameters. *IEEE Computer Graphics and Applications*, 10(6):44–55, 1990.

[32] Lixing Yu, Atsutake Kosuge, Mototsugu Hamada, and Tadahiro Kuroda. An anomaly detection system for transparent objects using polarized-image fusion technique. In *2022 IEEE Sensors Applications Symposium (SAS)*, pages 1–6, 2022.

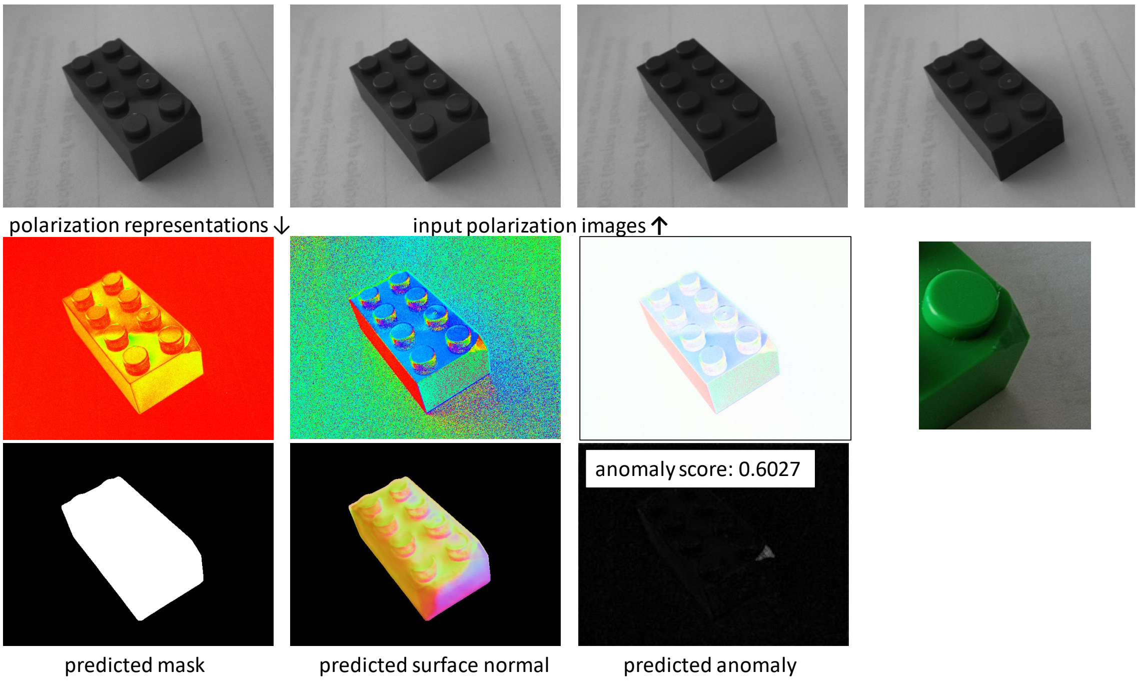

(a)

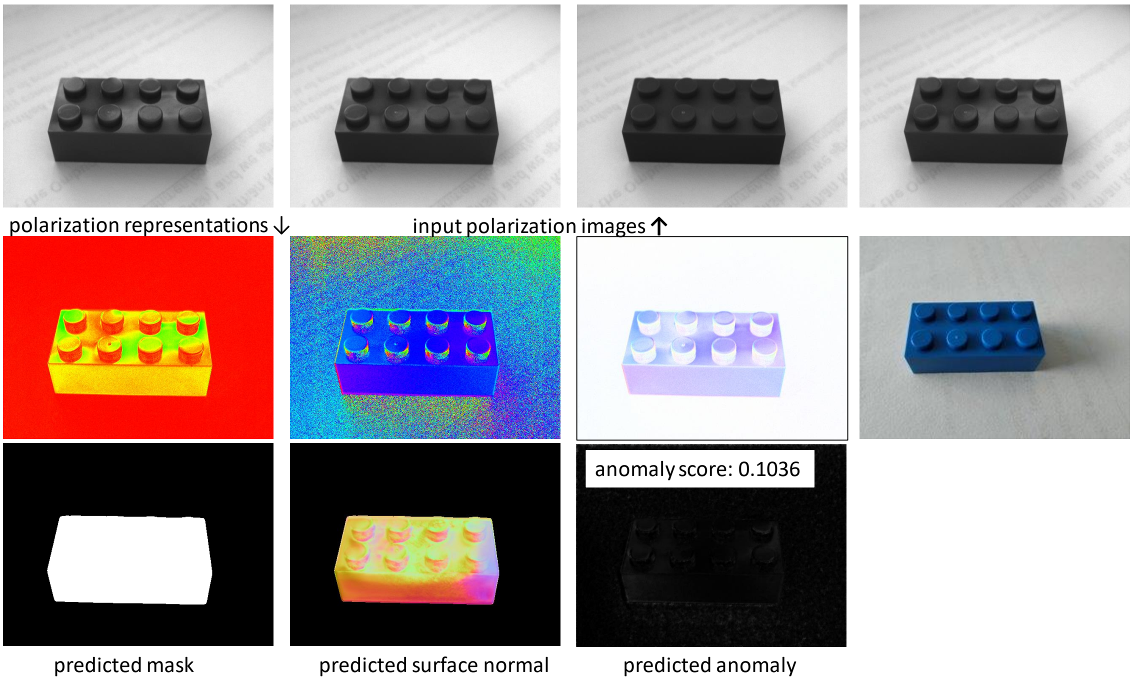

(b)

Fig. 4: **Visualized Results of Surface Anomaly Detection for Real Data Sample**: (a) sample with structural damage on the corner, and (b) the defect-free sample.

[33] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8330–8339, 2021.