International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
IJISAE
ISSN:2147-6799 www.ijisae.org Original Research Paper

Self-Supervised Learning Methods for Limited Labelled Data in Manufacturing Quality Control

[1]Suresh Sankara Palli, [2]Niranjan Reddy Rachamala, [3]Sukesh Reddy Kotha, [4]Manasa Talluri

Submitted:02/11/2021 **Revised:** 18/12/2021 **Accepted:** 25/12/2021

Abstract: Successful deep learning model deployment often depends on the quantity, quality, and accessibility of annotated data, as the use of deep learning methods in industrial applications expands at an accelerating rate and scale. The issues of effective data labelling and annotation verification in a human-in-the-loop scenario are addressed in this work. The Laser-based Directed Energy Deposition (L-DED) procedure is the subject of this work, which makes use of embedded vision systems to record crucial melt pool properties for ongoing observation. In order to provide in-situ monitoring without ground truth information, two self-learning frameworks based on Transformer architecture and Convolutional Neural Networks are deployed to analyse zone pictures from various DED process regimes. Although they need explicit human supervision, deep convolutional neural networks have recently shown respectable improvement in learning spatial patterns in WBMs. Furthermore, the RGB pictures that make up the majority of these datasets vary greatly from X-ray images. To overcome this drawback, our study suggests an approach that uses X-ray imaging and domain-specific self-supervised pretraining methods to enhance the ability to identify defects in manufactured goods. To improve feature extraction from manufacturing photos, we use SimSiam and SimMIM, two pretraining techniques. An industrial dataset of 27,901 unlabelled X-ray pictures from a manufacturing production line is used for the pretraining phase. Furthermore, we highlight how the models pretrained using X-ray pictures have improved their capacity to identify important flaws, which is essential for maintaining safety in industrial environments. Significant proof of the advantages of self-supervised learning in manufacturing defect identification is provided by our study, laying the groundwork for future investigations and useful applications in industrial quality control.

***Keywords: -** Laser-based Directed Energy Deposition (L-DED), X-Ray Images, Solid Foundation, Manufacturing Products, Convolutional Neural Networks, Industrial Dataset, Deep Learning Techniques.*

I. Introduction

Many studies have been carried out over the years to improve the visual examination of industrial goods via the use of X-ray imaging [1]. The development of automated procedures that can detect faulty items has been the primary emphasis. In reality, operators find that manually analysing every component is not only a tiresome and repetitive operation, but it also tends to reduce their accuracy with time [1, 2]. On the other hand, data-driven methods successfully reduce the possibility of human mistake while also guaranteeing consistent performance over extended periods of time. They may thus greatly facilitate the operators' decision-making process.

[1]*Independent Researcher, USA.*

[2]*Independent Researcher, USA.*

[3]*Independent Researcher, USA.*

[4]*Independent Researcher, USA.*

Recent developments in deep learning-based methods have made them the best option for a wide range of jobs in many different fields. In particular, these technologies have greatly outperformed conventional approaches and are currently regarded as state of the art in the field of manufacturing fault identification. However, their data-hungry character is a crucial condition that affects their performance [2, 3]. For these methods to develop efficient visual representations, large datasets of labelled pictures are needed during training. As a result, their efficacy may be significantly reduced when a small number of photos are available, indicating a significant problem in their use [3, 4]. This difficulty is made much more apparent in defect identification since it is particularly challenging to get large, precisely labelled information in industrial settings.

Existing research has shown effectiveness across a wide range of domains and tasks, but at the expense of an increasing quantity of data and computation, deeper learning algorithms and models continues to advance with increasing capacity & complexity.

Nevertheless, a large number of commercial applications lack easily accessible, high-quality datasets [3, 4]. Consequently, a significant portion of the machine learning life cycle involves data engineering, which often calls for laborious, costly, and time-consuming human annotation and inspection.

Automating the data curation processes and lowering the number of labels required for optimal performance are essential to minimise the amount of human labour [6, 4]. By selecting the most informative sample of data for labelling, for instance, active learning may reduce the amount of human labour needed [4, 5]. Recent advancements in active learning have shown encouraging outcomes in terms of accelerating annotation by humans in the loop.

By helping models learn with fewer labels, both self-supervised learning and semi-supervised training have achieved competitive results compared to supervised baselines with less supervision [5, 6]. Motivated by recent advancements in semi-supervised and self-supervised learning techniques, we combine the finest aspects of both methodologies to provide a simple yet flexible framework for accurate and efficient data verification and image similarity-based classification.

DED is a very flexible additive manufacturing process that allows for accurate material deposition and has many uses in a variety of sectors. The fabrication of near-net-shaped components, feature enhancement, and repair are the three main types of DED applications. DED has significant benefits over traditional manufacturing methods because it reduces material waste by producing near-net-shaped components [2, 9]. Additionally, this technique makes it easier to fabricate novel, fairly bulky components with location-dependent features and little to no tooling needs. Additionally, it is useful for adding material to existing components to improve their performance.

Through DED, where more material is accumulated, worn-out or damaged portions in components may be effectively repaired and renovated [6, 9]. DED's adaptability is further shown by its capacity to blend disparate metals to produce functionally graded structures with different material qualities. DED's capacity to handle novel materials and multiple materials with remarkable efficiency is one of its main advantages.

For a variety of uses, including fast prototyping, tool or die repair, coatings, surface modification, and even the manufacturing of large-scale components [6, 9], DED has been extensively embraced by industries including aerospace, automotive, machining, and medical. Interestingly, DED may be used with a variety of materials, such as metals, alloys, and composite. In DED, a metal wire or powder is inserted at the centre of a thermal energy source, where it melts and fuses to the substrate or layers that have already been produced. Depending on the material that is being processed, the intended use, and the particular DED technology being used, a variety of heat sources may be used in DED procedures [6, 9]. Lasers, plasma, and electric arcs are the main heat sources employed in DED. More widely used, laser-based DED (L-DED) provides excellent precision and precise control, making it appropriate for a variety of materials and applications.

Since the melt pool is used for deposition and is correlated with all of the basic parameters that control the L-DED process, it is logical to assume that its condition will have a major impact on the produced part's quality [7, 8] Fig. 1. [9] provides a didactic example of the two methods utilised in DED to collect emissions from the process zone, or melt pool: co-axial and off-axial process zone sensing. The present condition of the DED process may be determined by analysing the emissions from the melt pool.

The sensing and monitoring apparatus is positioned on the same axis as the energy beams utilised for material deposition in co-axial process monitoring. In order to monitor the process in real time and provide prompt feedback, fault identification, and control changes, co-axial process monitoring combines sensors and cameras with the energy beam [9, 10]. However, in order to record certain process attributes without interacting with the energy source, off-axial process monitoring positions sensors away from the deposition zone. These sensors may be positioned to record certain elements of the operation, including the heat-affected zone or the molten material plume [11]. Additionally, they may watch the development of the deposited material or keep an eye on the workpiece's temperature distribution.

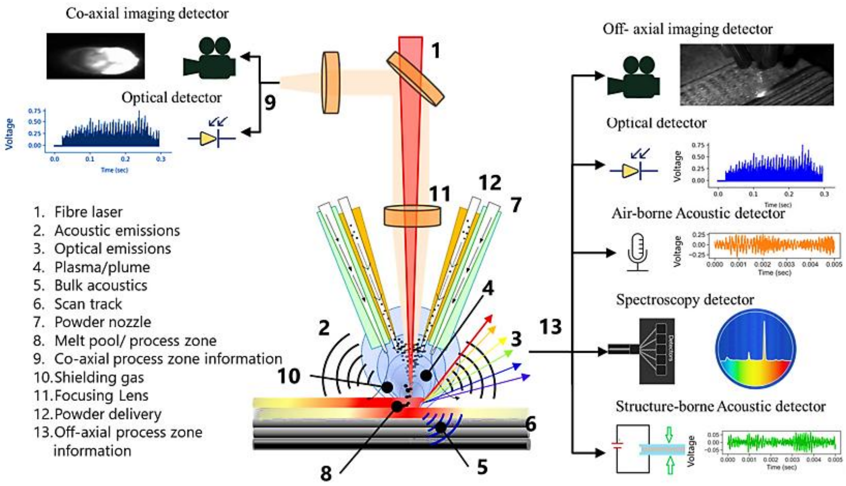

Fig. 1 Co-axial and off-axial process zone sensing are two approaches used in DED (Color figure online). [12]

Improved decision-making and increased process efficiency result from the gathering and decoding of vital data made possible by the use of non-intrusive sensors in the process zone in conjunction with sophisticated algorithms [14]. Furthermore, the environment of the process is not changed by these non-intrusive sensors, enabling undisturbed observations. Their capacity to monitor continuously [12,13] is a major benefit of such systems, outperforming post-mortem inspection methods and intermittent machine diagnostics.

If low component quality is identified, this real-time monitoring enables prompt action, averting more irregularities and providing improved production process management. In the end, these developments increase DED's accuracy, dependability, and quality, increasing its worth as a production technique across a range of sectors [14]. The list of non-intrusive sensors that have been reported in the literature and are capable of recording various process zone characteristics from laser-based DED and distinguishing between steady state and abnormality is shown in Table 1. The majority of events in the process zone are temporary because laser-based processes entail complex physics [15]. The multidimensional data from sensors that may record such occurrences must be analysed smoothly, and judgements must be taken concurrently with little human involvement, in order to guarantee that online diagnostics is a sensible substitute for quantification.

The ability to identify nonlinear patterns in data, learn from them, and eventually make judgements is made possible by the Machine Learning (ML) paradigm and soft computing approaches [16]. The advantages of combining sensing techniques with machine learning to monitor the DED process have been shown in earlier studies.

The traditional method of depending on skilled engineers to visually examine WBMs is no longer acceptable and has to be replaced with a revolutionary approach. It is almost hard for engineers to manually navigate through every WBM and assign them to one of the predetermined fault classifications since contemporary fabs manufacture thousands of wafers each week [16]. Depending on their experience levels or working situations, engineers often make biassed conclusions.

Additionally, it is possible for previously unidentified fault patterns to surface early on in the process development of new production stages [11]. Even while WBM data streaming during large-scale manufacturing may include valuable information for identifying the underlying cause and so producing improvements, an almost incalculable amount of this data is left unprocessed [19]. The examination of WBM failure patterns will undoubtedly benefit from automated methods that can take use of this enormous volume of unanalysed data. Unsupervised learning is a potential approach for creating fully automatic classifiers in a situation where labelled data is limited and unlabelled data is plentiful. Recent work in computer vision [10] has shown the ability to learn high-quality feature representations from large-scale unlabelled pictures alone. This study falls under the general heading of self-supervised learning.

A hybrid clustering technique that uses a spherical shell algorithm and Gaussian expectation-maximization-based spatial filters to find fault patterns [12]. Defective WBMs have spatial autocorrelations, and dynamic warping is used for grouping and spatial correlograms for the extraction of features. Self-organised maps were trained to find typical defect patterns, and SVMs were trained using the given clusters as hard labels. Using a morphology-based SVM to create synthetic samples for WBM similarity searches [19].

When adequate class labels are not available, self-supervised learning in deep learning seeks to create strong feature representations straight from the input data. Creating supplementary pretext tasks that motivate a neural network to acquire useful representations is the main issue of self-supervised learning [12,]. This new learning paradigm has extended out to a variety of data fields, including artificial intelligence, signal processing, processing of natural languages, or robotics, by taking use of the high-level semantics of the incoming data [, 21].

Our focus is limited to the picture domain. Predictive techniques, in which one portion of the input picture is excluded and can be predicted given the other portions of the image, are the focus of one line of study in that field. Predicting the relative positions of patches cropped from a single image [12], matching exemplars [13], inpainting patches that are absent based on the remaining context , resolving jigsaw puzzles [19], colourizing images from greyscale to RGB [21], and anticipating image orientations [23] are just a few examples of research in this area.

These datasets should include events and objects that can be uniquely identified by state and class thanks to their properties. Some datasets are easier to access, such those that include recordings of everyday tasks like cooking [1], [2]. However, datasets may be small in size or lack the thorough annotations needed for the creation of AI algorithms for specialised tasks, such the identification of intricate technical processes. Nevertheless, data sets may not be accessible in sufficient quantity or with the required annotations for the creation of AI algorithms for some specialised tasks [10], such as the identification of technical processes. Consequently, more data must be gathered, important steps must be highlighted, and pertinent annotations must be made in order to analyse certain processes.

The second significant obstacle to the effective use of AI in manufacturing & operational process monitoring systems is the need for the creation of specialised, highly accurate algorithms. Such algorithms must be able to learn from tiny annotated data sets in order to provide flexible solutions that can adjust to an increasing number of technological processes. Self-supervised learning is a useful strategy that involves pre-training a model on a large amount of unlabelled data before refining it on a smaller amount of labelled data [19]. The model learns to recognise fundamental patterns in the data during the first phase, known as the pretext task. At this point, the tasks and process types change. Using task-relevant information and modifications to its final layers, the model is customised for the particular application in the second step, known as the downstream task.

The use of transfer learning is one possible way to lessen the data needs. This strategy essentially consists of transferring information gained from challenges with broader definitions to activities that are more focused and particular [19]. Pretraining vision models on large datasets such as ImageNet improves their performance on a range of downstream assignments [19]. But even though these models have been shown to be successful in detecting manufacturing defects, their performance may be affected by (1) the discrepancy between ImageNet and X-ray pictures and (2) their possible bias towards the particular categories that are present in the dataset.

The creation of deep learning models is often hampered by these pretraining methods' need for a labelled dataset. A method that makes effective use of the enormous amounts of data that are accessible is self-supervised learning, which makes it possible to extract important characteristics from photos without depending on labelled data [14, 19]. Self-supervised techniques enable the use of large datasets of unlabelled X-ray pictures, which are typical in industrial settings, in the context of manufacturing fault identification. This makes it easier to train a model that can reliably extract the most relevant characteristics and is skilled at comprehending X-ray picture representations. The model then serves as a good foundation for further tasks using a tagged X-ray dataset, increasing its usefulness and efficacy in accurately identifying defects.

Over the years, a lot of work has been done to use computer vision techniques to automate the process of detecting manufacturing defects. Nowadays, detecting manufacturing defects is a well-known computer vision issue that is used extensively in many industrial quality control procedures . Prior to using statistical methods for feature extraction, early efforts at defect identification included picture comparison and the Fourier Transform to find flaws. Machine learning was then used to classify these attributes in order to differentiate between objects that were faulty and those that weren't. However, this laborious feature extraction limited the transferability of acquired information since it was product particular and not generally applicable.

As was already noted, it is well known that in order for modern deep learning algorithms to acquire intrinsic representations of data and achieve enough generalisation ability, they need large-scale training datasets. Large datasets must be annotated since supervised learning requires tagged data for training [23]. The lengthy and costly labelling procedure, which is also impractical in a number of industries, might be seen as a bottleneck. Additionally, supervised models rely significantly on human annotated labels.

II. Methodology

To improve the existing state of the art of manufacturing defect detection, our approach is based on the use of self-supervised learning methods to build defect detector for manufacturing [2, 3]. By using large datasets of unlabelled X-ray images—which are often seen in industrial settings—to train defect detectors in a self-supervised way, it differs from earlier methods in manufacturing defect identification.

Our approach's pretraining stage entails creating models that can automatically recognise and extract important visual characteristics from X-ray pictures without the need of annotated data. At this step, an unlabelled database of X-ray manufactured pictures is used to train the models in a self-supervised manner [23]. We describe in detail two different self-supervised learning methods in the next parts of this study [26]. These techniques were especially picked because of their proven effectiveness and high results on previous pretraining exercises. We take use of the natural features and patterns seen in these photos as a consequence of the training process, which allows the models to detect minute but important irregularities that indicate flaws. Furthermore, there are no significant expenses associated with data collection during this pretraining phase.

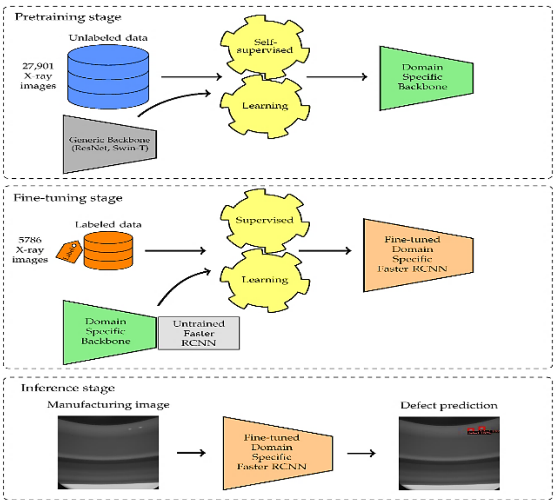

Fig. 2 Shows an overview of the methodology, which encompasses two key stages: pretraining and fine-tuning. In the pretraining phase, the backbone of the model is trained through a self-supervised learning approach [19].

We used a feature extractor that was specially designed to distinguish between relevant characteristics from manufacturing photographs for pretraining. SimSiam was selected for picture classification tasks because it offers the best balance between ease of use and efficacy [23]. Its benefit is that it can be trained on numerous GPUs without using a lot of resources since it doesn't need big batch sizes or an automatic momentum encoder.

$$\mathfrak{D}(p_{1,}Z_2) = -\frac{p_1}{||p_1||_2} \cdot \frac{z_2}{||z_2||_2}$$

$$\mathfrak{D}(p_{1,}Z_2) = -\frac{p_2}{||p_2||_2} . z\frac{z_1}{||z_1||_2}$$

$$\mathcal{L} = \frac{1}{2}\mathfrak{D}\left(p_{1,}stopgrad(z_2)\right) + \frac{1}{2}\mathfrak{D}\left(p_{2,}stopgrad(z_1)\right)$$

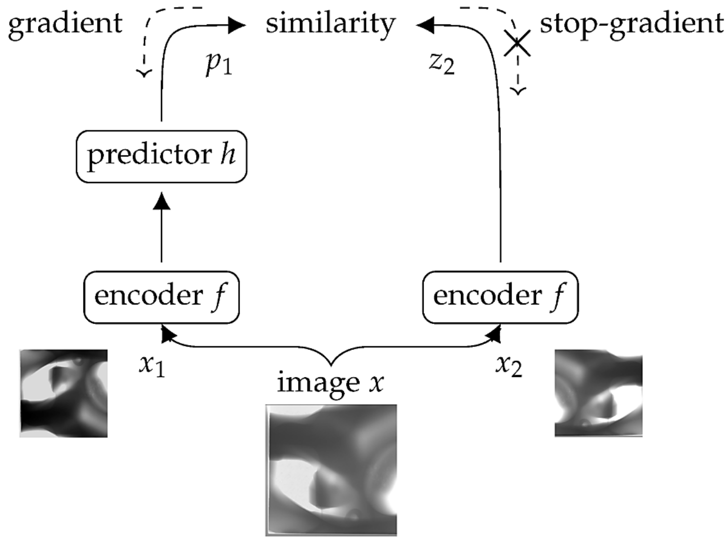

Fig. 3 Overview of the SimSiam architecture. [11]

Transformers have become the state of the art for many vision-related tasks, mainly because of their remarkable representation learning and feature extraction capabilities [19]. Our investigation into their use for manufacturing process flaw identification was motivated by this advancement.

$$\mathcal{L} = \frac{1}{\Omega(x_M)}||y_M - x_M||_1$$

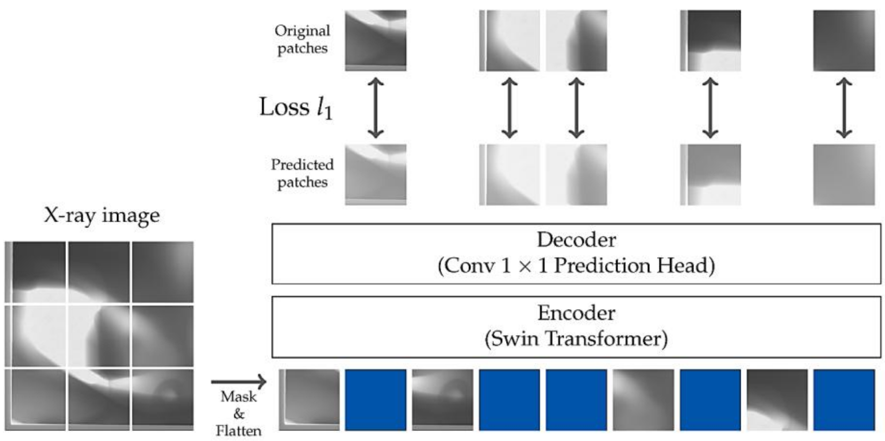

Fig. 4 Overview of the SimMIM architecture. [19]

III. Implementation And Results

Here, we describe the tests that were done to see how well the models worked after being pretrained on X-ray pictures of manufacturing items [12]. We describe the methods used and the particular datasets used in our research. The findings are also discussed in this part, providing information on how various pretraining strategies affect performance [24].

We conducted a model pretraining step on X-ray manufacturing pictures using the two self-supervised techniques described above, SimSiam and SimMIM [23]. We postulated that self-supervised preparing on a large collection of X-ray pictures may enhance models' performance in downstream tasks involving the identification of manufacturing defects.

Given that there are seven critical defect types and twelve minor defect types in the defect distribution, finding flaws in the product and categorising them into one of the two groups provide a difficult defect detection task. With 4002 samples for minor flaws and 1784 samples for major defects, there was a noticeable class imbalance in the dataset distribution . Table 1 describes how the target dataset was divided into training and test sets while preserving class proportions.

Fig. 1 Distribution of images for critical and minor defects the target dataset. [10]

	Critical defects	Minor defects
Train	1423	3200
Test	369	809

First, we used SimSiam and SimMIM to train a backbone model. In both situations, we trained each model for 100 epochs using the default approach from the original articles. The normal assessment technique for the models that were previously trained was then put into practice; [19] we added a linear layer to the top of the model while freezing the other backbone layers. We used the target dataset with both minor and significant flaws to train our classification classifier. A model using ImageNet pretrained weights was used to create a baseline. The feature extractor we utilised for SimSiam was ResNet50, which produced a 2048-dimensional structure feature vector for our classification layer. On the other hand, we used the original Swin Transformer that was set up with a 12-pixel window size for SimMIM [10, 19]. In the Swin Transformer example, the feature vector has 1024 dimensions [25]. Table 2 provides specifics of this linear classification's outcomes on our target dataset.

Table 2 Linear classification results on industrial target dataset for both RestNet and Swin-T Backbones. [19]

Backbone	Weights	AP	Accuracy
ResNet	ImageNet	49.6	69.6
	X-ray SimSim	79.5	72.6
Swin-T	ImageNet	89.6	83.9
	X-ray SimMiM	78.9	69.9

We used three sets of initial weights to train the Faster R-CNN model on the industrial target database in the first step of this fine-tuning phase: those derived via X-ray pretraining using both the SimSiam and SimMIM approaches, as well as those initialised using ImageNet. This approach was used to assess how well the models handled challenging, real-world defects identification tasks in the industrial sector [19]. According to Table 3, the results show that the SimSiam and SimMIM pretraining techniques outperform the traditional supervised methodology, particularly when it comes to detecting important faults [10].

Table 3 Fine-Tuning results on defect detection for manufacturing dataset.

Dataset	Backbone	Pretraining	(m) AP	AP-critical
Industrial	ResNet	ImageNet	89.6	94.6
		X-ray SimSim	87.9	93.7
	Swin-T	ImageNet	91.6	96.9

		X-ray SimMiM	92.6	**96.8**
	ResNet	ImageNet	96.9	
		X-ray SimSim	94.8	
	Swin-T	ImageNet	93.0	
		X-ray SimMiM	99.5	

To illustrate the effectiveness of our detection of defects approach, we instead concentrate on the GDXray dataset. In particular, we use the SimSiam method to demonstrate the efficacy of the ResNet backbone previously trained using X-ray images [19]. We hope that Figure 4's visualisation will shed light on the usefulness and practical implementation of our methodology in spotting flaws in industrial settings.

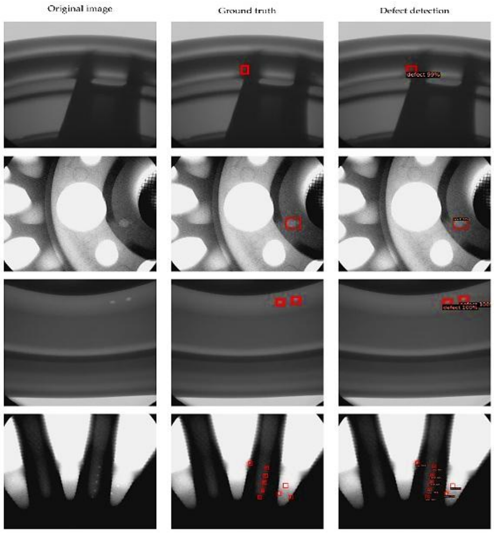

Fig. 5 Results of defect detection on the GDXray dataset. This figure illustrates the original X-ray image, ground-truth annotations, and the predictions of the model.

IV. Discussion

These tests provide insight into how well X-ray image-based pretraining techniques work, particularly when it comes to identifying manufacturing flaws. Our tests demonstrate that when it comes to identifying pertinent features for defect identification, models pretrained using X-ray pictures often outperform those pretrained with the ImageNet [21]. This domain-specific pretraining's advantage emphasises how crucial it is to match the pretraining stage with the distinct qualities of the task-specific visuals.

When there is a substantial amount of labelled data available, the Swin Transformer models outperform CNNs in comparing several backbones. The GDXray dataset, on the other hand, had fewer annotated pictures; in this case, the CNN backbones performed better. This implies that CNNs may work better with less data, even while Transformers perform better in settings with a lot of data [19].

The notable improvement in identifying essential abnormalities after pretraining with X-ray pictures is an important part of our study [23, 24]. This enhancement is particularly relevant in industrial environments where dependability and safety depend on precisely detecting such flaws. Significant ramifications for industrial applications may result from this improvement in defect identification, especially for important flaws [25]. Furthermore, pretraining on X-ray pictures showed better outcomes even with the smaller GDXray a database, demonstrating the models' flexibility in handling varying amounts of data. Given that data availability may fluctuate in real-world deployment scenarios, this flexibility is an essential quality.

V. Conclusion

This study adds to the continuous attempts to improve X-ray image-based flaw identification in industrial items. Our thorough tests and analyses provide a number of important conclusions that highlight the efficiency and usefulness of our suggested technique in real-world industrial situations. Our analysis demonstrates unequivocally that models previously trained on X-ray pictures perform better than those pretrained on ImageNet weights.

This observation is significant because it emphasises how domain-specific pretraining improves the model's capacity to identify relevant features for fault identification. We were able to get better detection skills than the models that were pretrained on more generic photos by matching the pretraining phase with the distinct features of the task-specific images. Additionally, comparing various backbone designs showed insightful patterns.

We found that CNN backbones perform better in datasets with less annotated pictures, whereas Swin Transformer models provide superior performance in circumstances with a large amount of labelled data. This suggests that based on the quantity of data provided and the particular needs of the activity, the best model may be chosen. The enhancement of critical defect identification, which is a key issue in industrial settings, made possible by pretraining on X-ray pictures is very remarkable. The innovations shown may greatly lower risks and improve dependability in industrial settings where accurate detection of such flaws is essential for guaranteeing safety and upholding high standards of quality.

Adoption of these improved detection systems may also significantly save maintenance costs and operational downtime, improve manufacturing processes, and boost industrial efficiency in general. Our approach helps to maintain the standing of industrial organisations and boost customer trust by improving product quality and reducing failure rates.

VI. References

[1] Wenbin Cai, Ya Zhang, Siyuan Zhou, Wenquan Wang, Chris Ding, and Xiao Gu. Active learning for support vector machines with maximum model change. In Machine Learning and Knowledge Discovery in Databases, pages 211–226. Springer, 2014.

[2] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. IJCV, 113(2):113–127, 2015.

[3] Yuhong Guo. Active instance sampling via matrix partition. In NIPS, pages 1–9, 2010.

[4] Suyog Dutt Jain and Kristen Grauman. Active image segmentation propagation. In CVPR, pages 2864– 2873, 2016.

[5] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In ICLR, 2018.

[6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In International conference on machine learning, pages 1597–1607. PMLR, 2020.

[7] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9729–9738, 2020.

[8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297, 2020.

[9] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 21271–21284. Curran Associates, Inc., 2020.

[10] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. arXiv preprint arXiv:2006.09882, 2020.

[11] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15750–15758, 2021.

[12] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder

networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 28. Curran Associates, Inc., 2015.

[13] Mery, D.; Arteta, C. Automatic Defect Recognition in X-ray Testing Using Computer Vision. In Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), Santa Rosa, CA, USA, 24–31 March 2017; pp. 1026–1035.

[14] Li, X.; Tso, S.K.; Guan, X.P.; Huang, Q. Improving Automatic Detection of Defects in Castings by Applying Wavelet Technique. IEEE Trans. Ind. Electron. 2006, 53, 1927–1934.

[15] Mery, D. Automated Radioscopic Inspection of Aluminum Die Castings. Mater. Eval. 2006, 65, 643–647.

[16] Tsai, D.M.; Huang, T.Y. Automated Surface Inspection for Statistical Textures. Image Vis. Comput. 2003, 21, 307–323.

[17] Zhao, X.; He, Z.; Zhang, S.; Liang, D. A Sparse-Representation-Based Robust Inspection System for Hidden Defects Classification in Casting Components. Neurocomputing 2015, 153, 1–10.

[18] Du, W.; Shen, H.; Fu, J.; Zhang, G.; He, Q. Approaches for Improvement of the X-ray Image Defect Detection of Automobile Casting Aluminium Parts Based on Deep Learning. NDT Int. 2019, 107, 102144.

[19] Mery, D. Aluminum Casting Inspection Using Deep Object Detection Methods and Simulated Ellipsoidal Defects. Mach. Vis. Appl. 2021, 32, 72.

[20] Kim, H., Cong, W., Zhang, H.-C., & Liu, Z. (2017). Laser engineered net shaping of nickel-based superalloy Inconel 718 powders onto AISI 4140 alloy steel substrates: Interface bond and fracture failure mechanism. Materials, 10(4), 341.

[21] Lei, J. B., Wang, Z., & Wang, Y. S. (2012). Measurement on temperature distribution of metal powder stream in laser fabricating. Applied Mechanics and Materials.

[22] Li, X., Siahpour, S., Lee, J., Wang, Y., & Shi, J. (2020). Deep learningbased intelligent process monitoring of directed energy deposition in additive manufacturing with thermal images. Procedia Manufacturing, 48, 643–649.

[23] Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., & Tang, J. (2021b). Self-supervised learning: Generative or contrastive. IEEE Transactions on Knowledge and Data Engineering.