Adapting Segment Anything Model (SAM) to Experimental Datasets via Fine-Tuning on GAN-based Simulation: A Case Study in Additive Manufacturing

Anonymous Author(s)
Affiliation
Address
email

Abstract

Industrial X-ray computed tomography (XCT) is a powerful tool for non-destructive characterization of materials and manufactured components. However, it faces significant challenges due to the complexity of internal structures, noise, and variability in resolution. Traditional computer vision models often struggle with noise, resolution variability, and complex internal structures, particularly in scientific imaging applications. State-of-the-art foundational models, like the *Segment Anything Model (SAM)*—designed for general-purpose image segmentation—have revolutionized image segmentation across various domains, yet their application in specialized fields like materials science remains under-explored. In this work, we explore the application and limitations of SAM for industrial X-ray CT inspection of additive manufacturing components. We demonstrate that while *SAM* shows promise, it struggles with out-of-distribution data, multiclass segmentation, and computational efficiency during fine-tuning. To address these issues, we propose a fine-tuning strategy utilizing parameter-efficient techniques, specifically *Conv-LoRa* , to adapt *SAM* for material-specific datasets. Additionally, we leverage generative adversarial network (GAN)-generated data to enhance the training process and improve the model's segmentation performance on complex X-ray CT data. Our experimental results highlight the importance of tailored segmentation models for accurate inspection, showing that fine-tuning *SAM* on domain-specific scientific imaging data significantly improves segmentation performance. However, despite improvements, the model's ability to generalize across diverse datasets remains limited, highlighting the need for further research into robust, scalable solutions for domain-specific segmentation tasks. Code and training data will be available in public.

No.	Size
Tr-1	5724
Tr-2	1142
Te-1	920
Te-2	920
Te-3	5000
Te-4	1080
Te-5	600
Te-6	785

Table 1: Dataset specification.

***SAM-GAN* Finetuning Settings:** We finetune on 4 GPUs of 80GB memory for finetuning. Learning rate set to .0003 and batch size was set to 4. Each training ran for 20 epochs. Most finetuning converge by 15 iterations. Each training finishes within 4 hours.
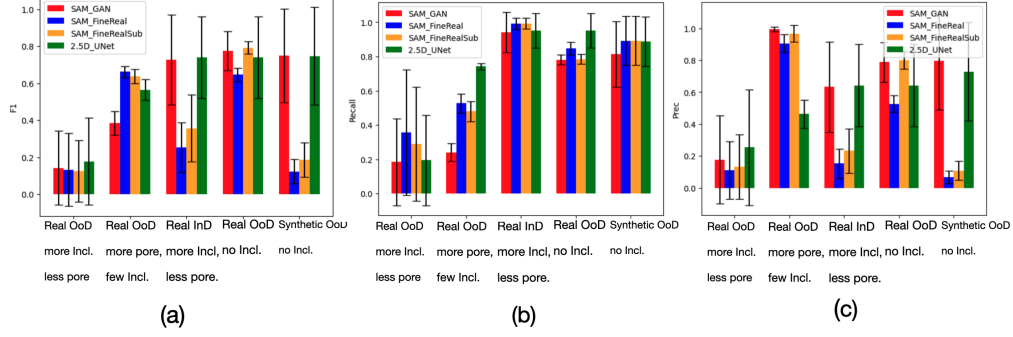
Figure 1: (a)-(c)(from left to right) mean-F1, mean-Recall, and mean-Precision values of *SAM* , *SAM-FineReal* , *SAM-FineReal-Sub* , and *U-Net* for the class *Pore*. It is noted that the recall values of *SAM* is lower for real OoD, while precisions are high in majortiy cases. This shows *SAM* identifies many as false negative pores than false positives in OoD.
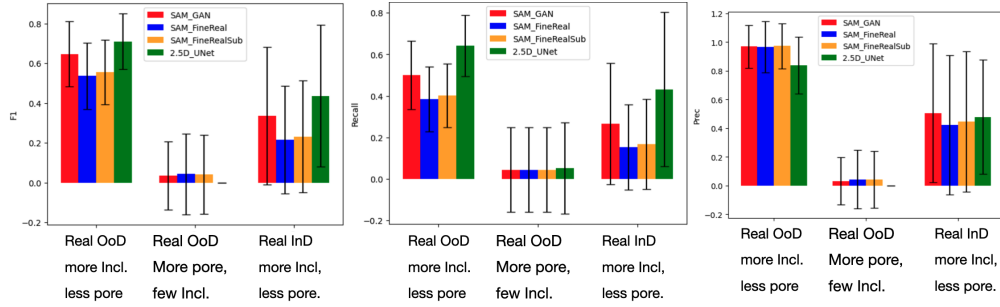

Figure 2: (a)-(c)(from left to right) mean-F1, mean-Recall, and mean-Precision values of *SAM* , *SAM-FineReal* , *SAM-FineReal-Sub* , and *U-Net* for the class *Inclusion*. It is noted that the recall values of *SAM* is lower for real OoD, while precisions are high in majortiy cases. This shows *SAM* identifies many as false negative pores than false positives in OoD.

0.1 Additional Results of OoD

Fig. 1 and Fig. 2 shows the mean-F1, mean-Recall, and mean-Precision for the OoD datasets (emphTe-1-*Te-5*) for class *Pore* and *Inclusion* respectively. FIg. 3 shows the performance using same metric on all classes for all test OoD data.
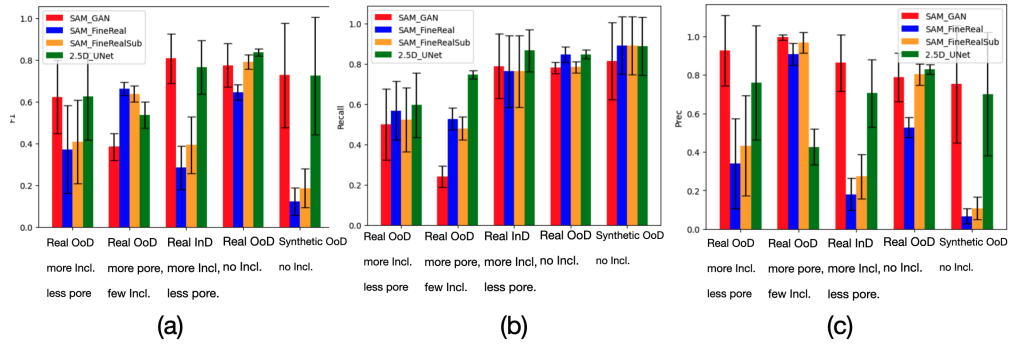
References

Figure 3: Analogous to Fig. 1, shows performance on all classes together. High value of F1 is dominated by high precision than recall.