UniDG

Paint the defect in. Leave the scene alone.

Unified Defect Generation via Defect-Context Editing

Yuanting Fan* Jun Liu* Bin-Bin Gao Xiaochen Chen
Yuhuan Lin Zhewei Dai Jiawei Zhan Chengjie Wang†

Tencent Youtu Lab NeurIPS 2026

One product does not teach the next

01

Real defects are rare

Inspection models learn from fault photos that a line seldom produces.

02

Few-shot generators overfit

A model fit to one defect category has to be trained again for a new object.

03

Editors miss the mask

A general editor changes the scene and loses the fine texture of the fault.

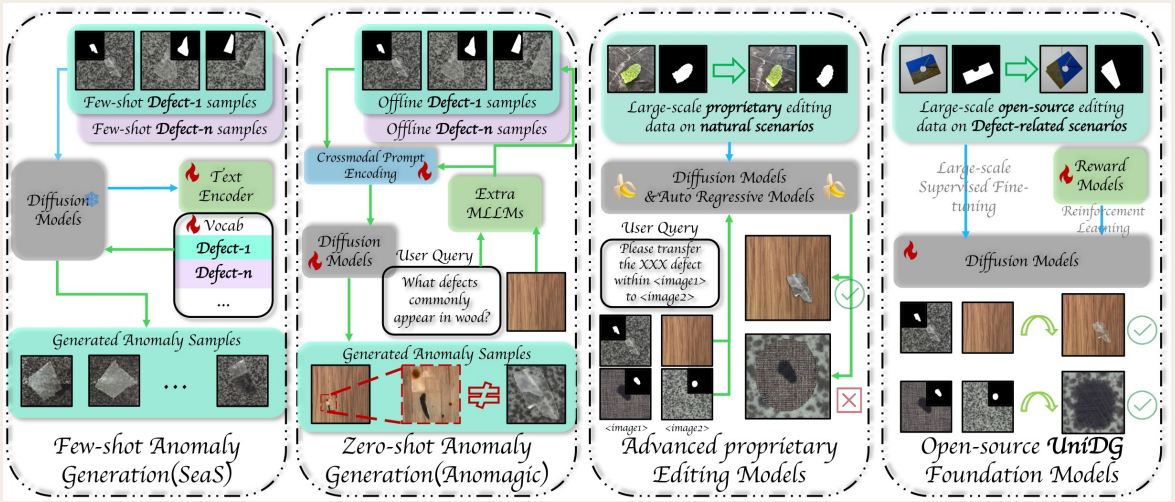
Keep the product. Edit only the defect.

A reference photo of the fault.

A new product, and a mask.

One model for every category.

The rest of the scene stays.



Train once, on paired edits

300K

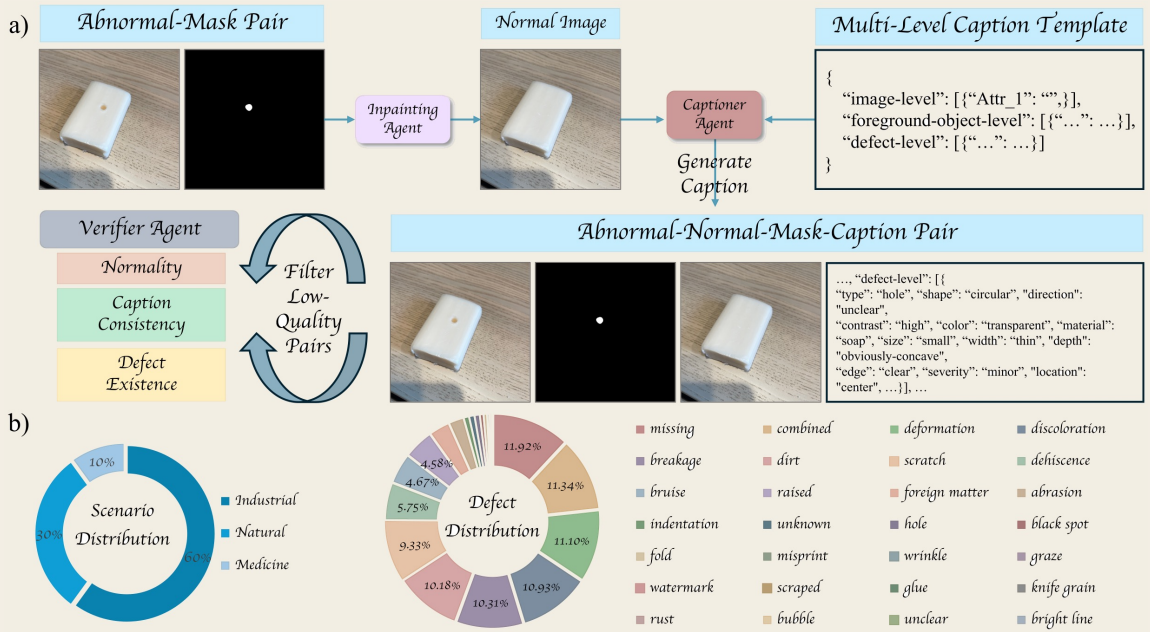
normal, defect, mask, caption

50

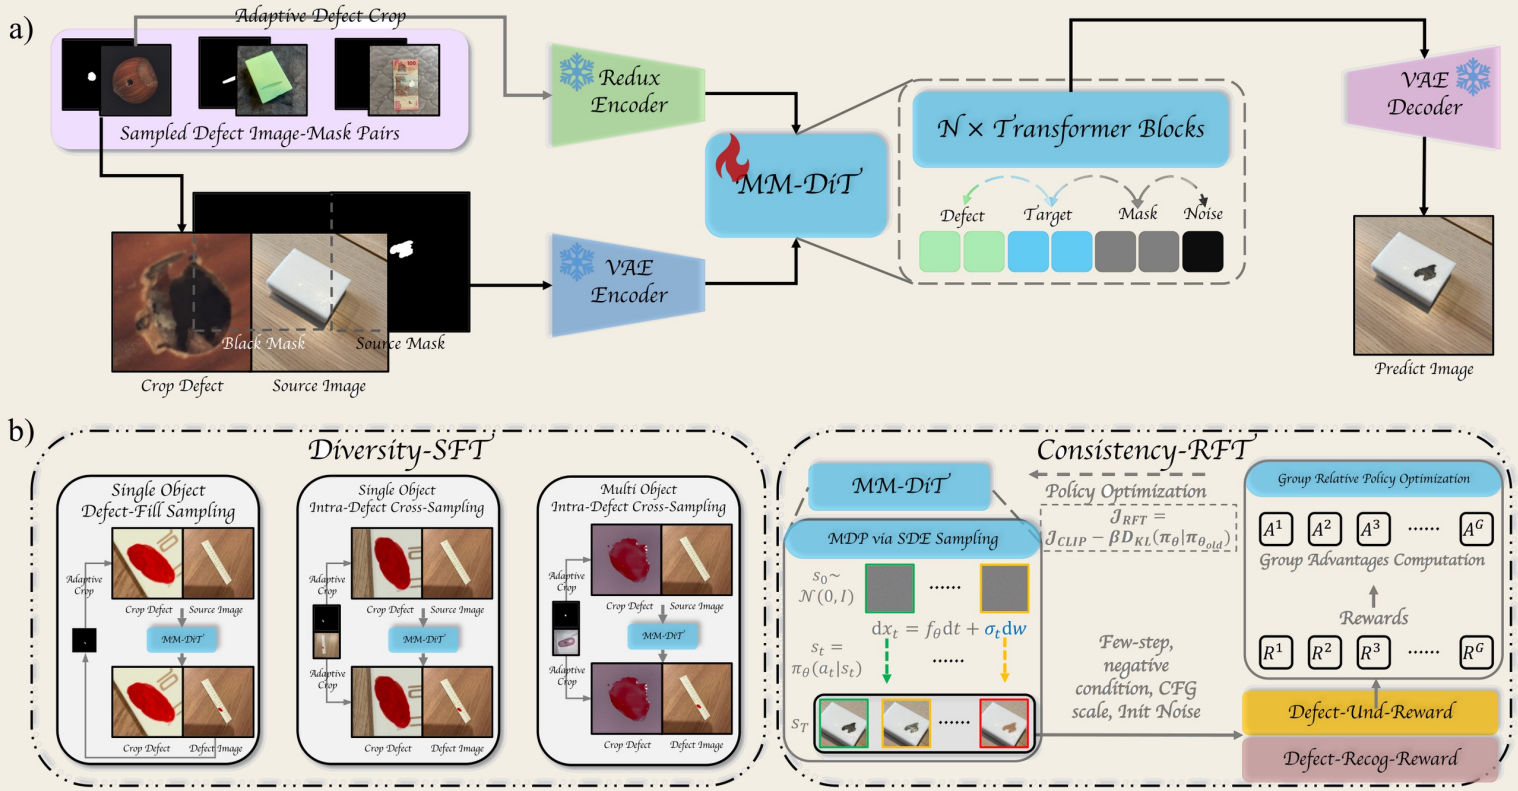
public datasets

Held out

MVTec-AD and VisA



Crop the reference. Edit on one canvas.



The cropped defect sits beside the masked product.

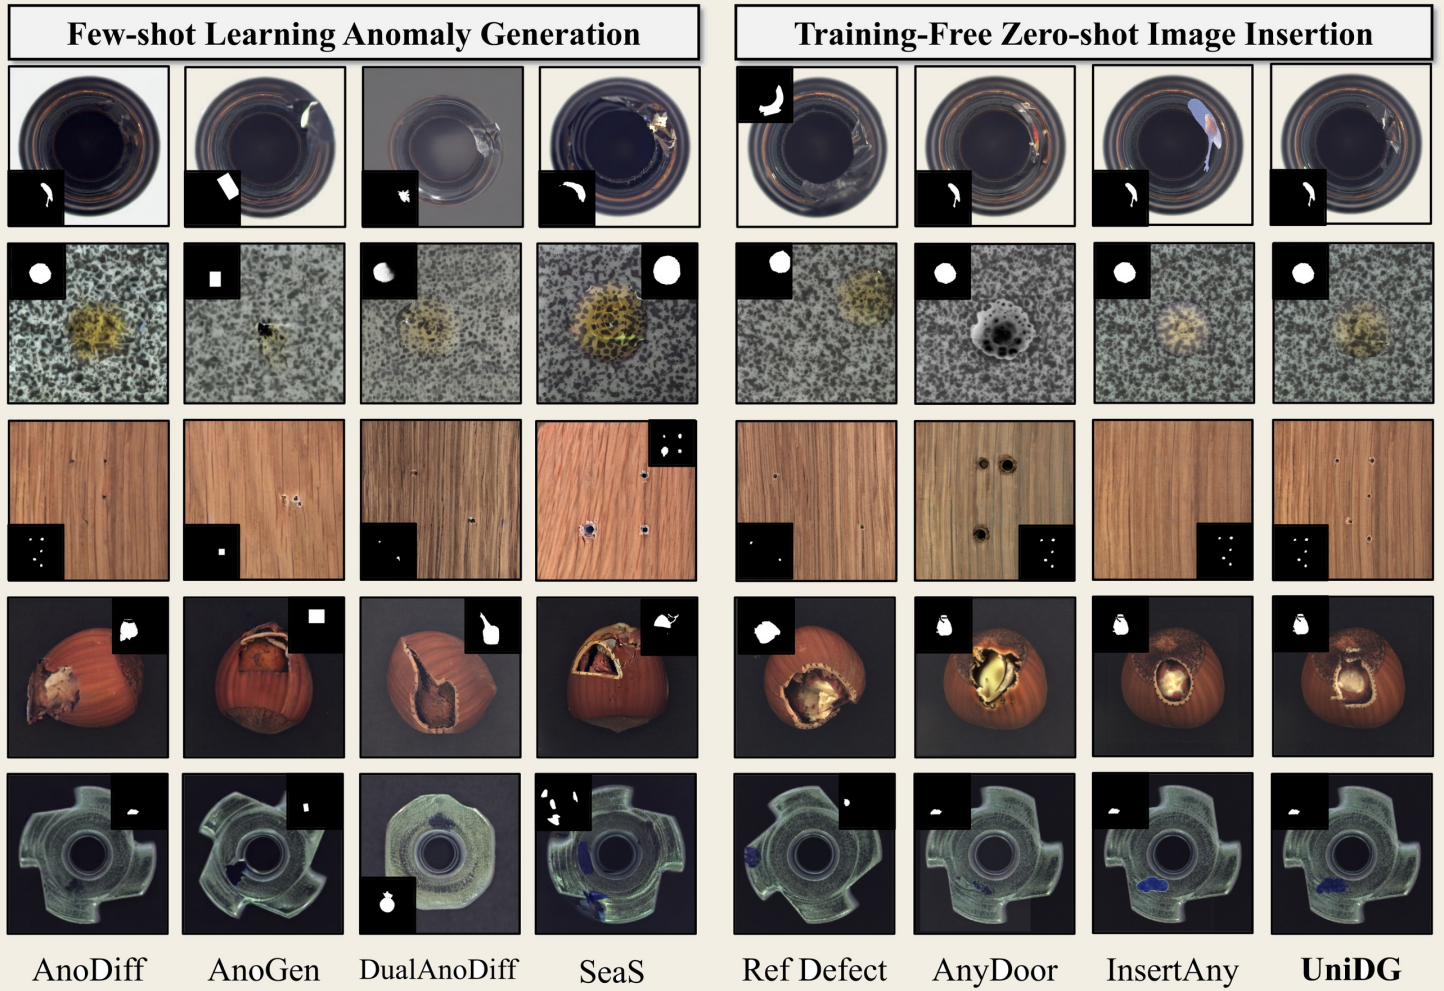
Diversity-SFT learns the transfer. A second stage can tighten the match.

Better defects.
Better detectors.

77.4 68.4
class accuracy

31.9 26.5
foreground IoU

91.1 81.8
region overlap



One model. Train once.

Leave the test products out.

PROJECT

retofan233.github.io/UniDG

PAPER

arXiv:2604.08915