Dissolving Is Amplifying: Towards Fine-Grained Anomaly Detection

Jian Shi[1] · Pengyi Zhang[2] · Ni Zhang[2] · Hakim Ghazahi[1] · Peter Wonka[1]

[1] KAUST · [2] NEC Laboratories China

Background:
Issues In Medical Anomaly Detection

The Long-Tail:

A few common conditions dominate while many rare conditions are underrepresented.

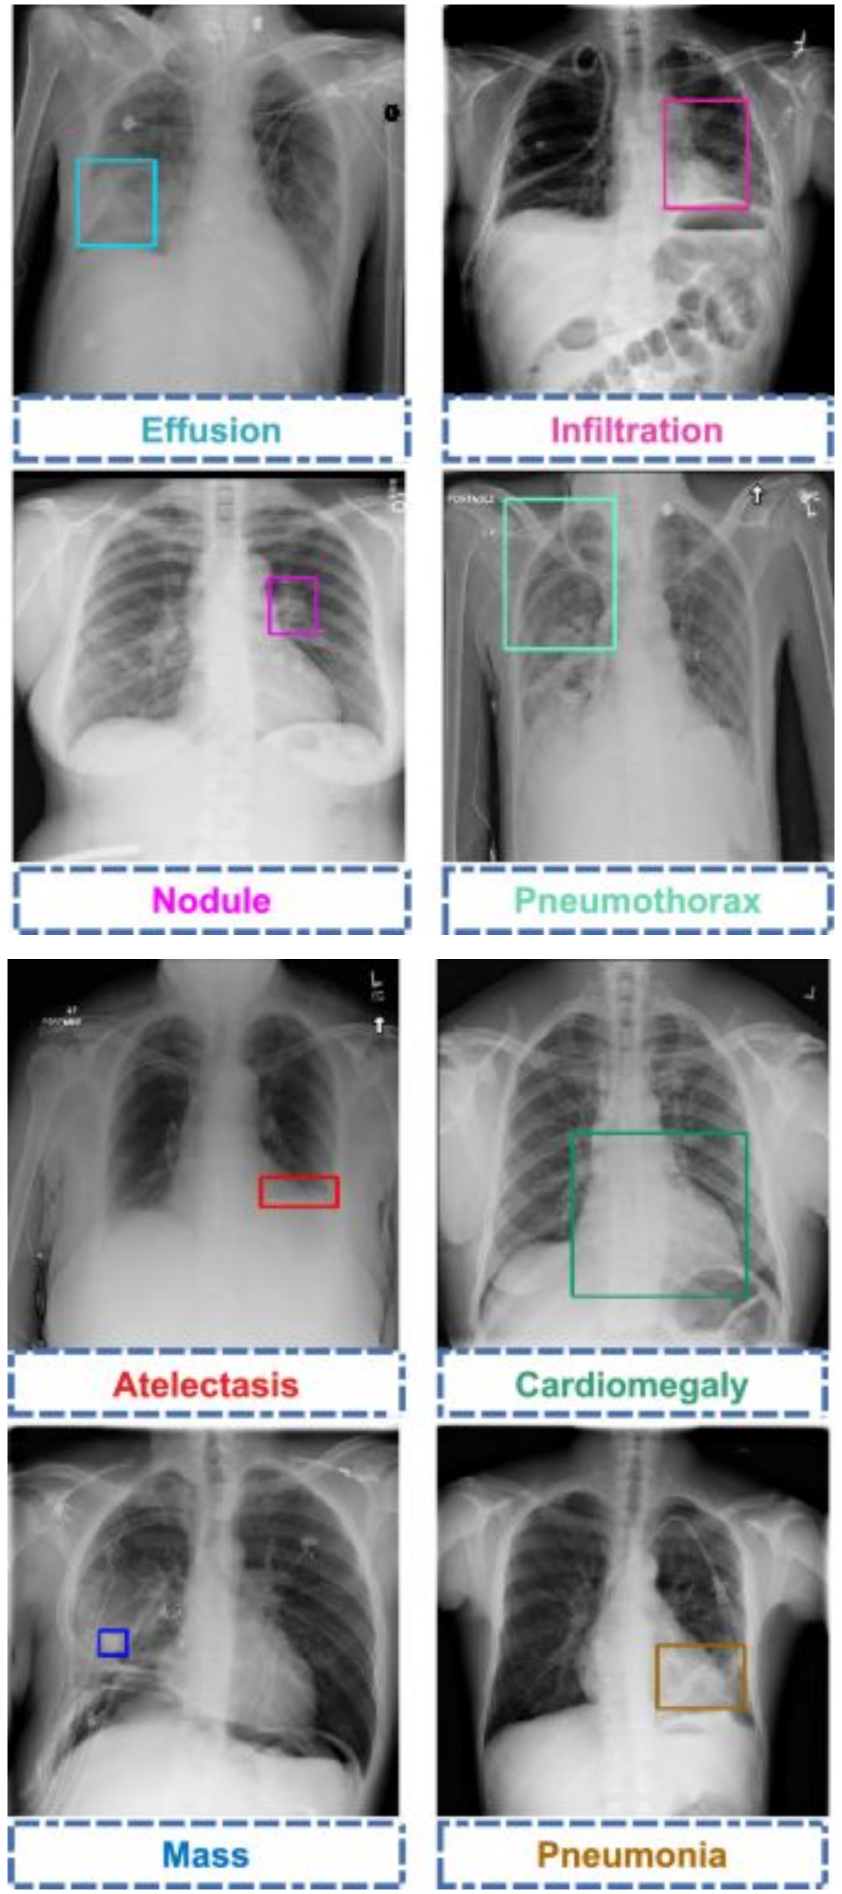

Healthy images tend to be quite similar, while unhealthy images can appear in a wide variety of forms.

Anomaly Detection in Medical Imaging

Dissolving Is Amplifying: Towards Fine-Grained Anomaly Detection

Challenge: How to focus on fine-grained details in medical images unsupervisedly?

Medical imaging often contains critical fine-grained features, such as tumors or hemorrhages, crucial for diagnosis yet potentially too subtle for detection with conventional methods

Dissolving Transformations w/ Diffusion Models

Employs dissolving transformations to detect subtle features, such as tumors, that traditional methods often overlook.

Amplifying Framework For Anomaly Detection

The DIA framework is designed to improve fine-grained anomaly detection in medical imaging, providing a robust solution to identify irregularities.

Significant AUC Improvement

Demonstrated a substantial AUC improvement of 18.40% over baseline methods across six diverse medical datasets.

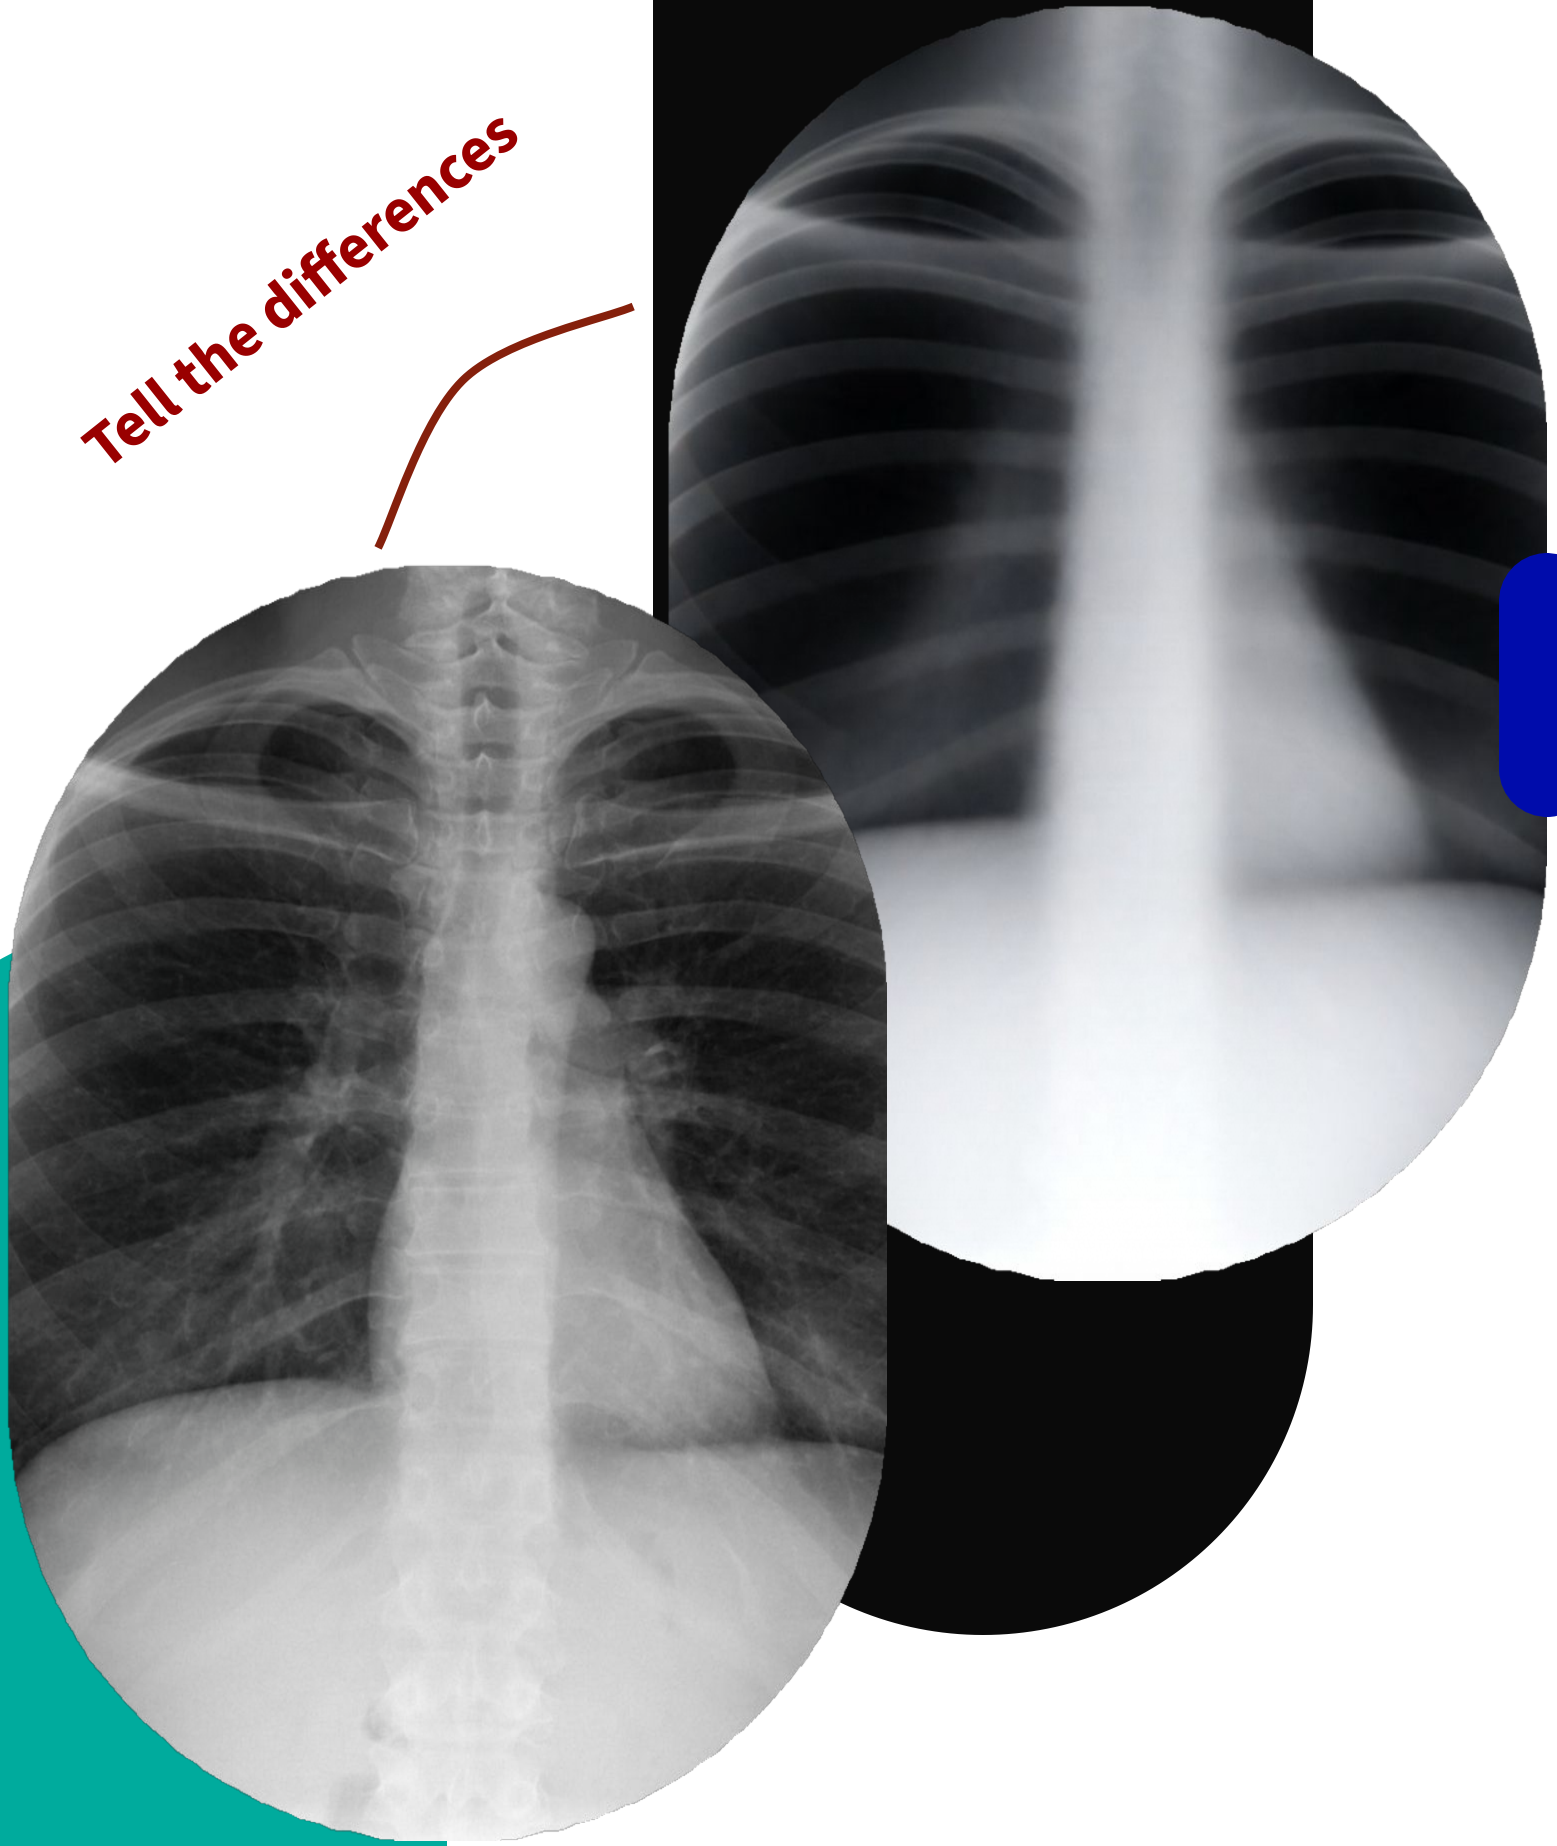

Our Approach:

Remove the fine-features from the image

&

Compare with the original image

&

Learn the differences

Key Technology:

Dissolving Transformations w/ Diffusion Models

Regular Diffusion:

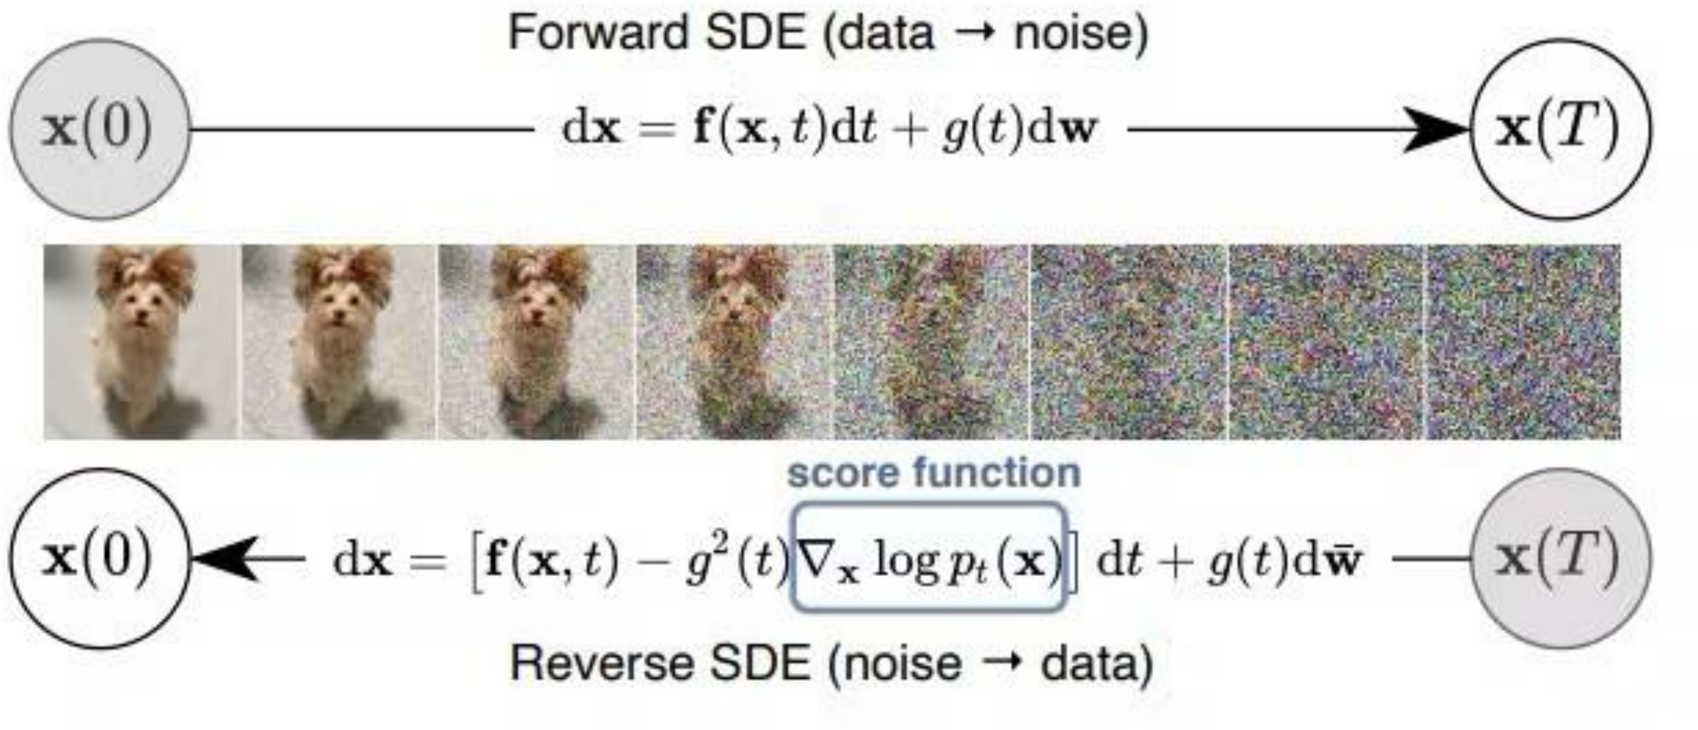

Dissolving Transformation:

Instead of generating images by progressive denoising, we apply reverse diffusion in **a single step**.

$$\hat{x}_{t\to 0} = \sqrt{\frac{1}{\bar{\alpha}_t}} \cdot x - \sqrt{\frac{1}{\bar{\alpha}_t} - 1} \cdot \epsilon_\theta(x, t),$$

$$\bar{\alpha}_t := \Pi_{s=1}^{t} \alpha_s \text{ and } \alpha_t := 1 - \beta_t,$$

Dissolving Transformation Effects

Instead of generating images by progressive denoising, we apply reverse diffusion in a single step.

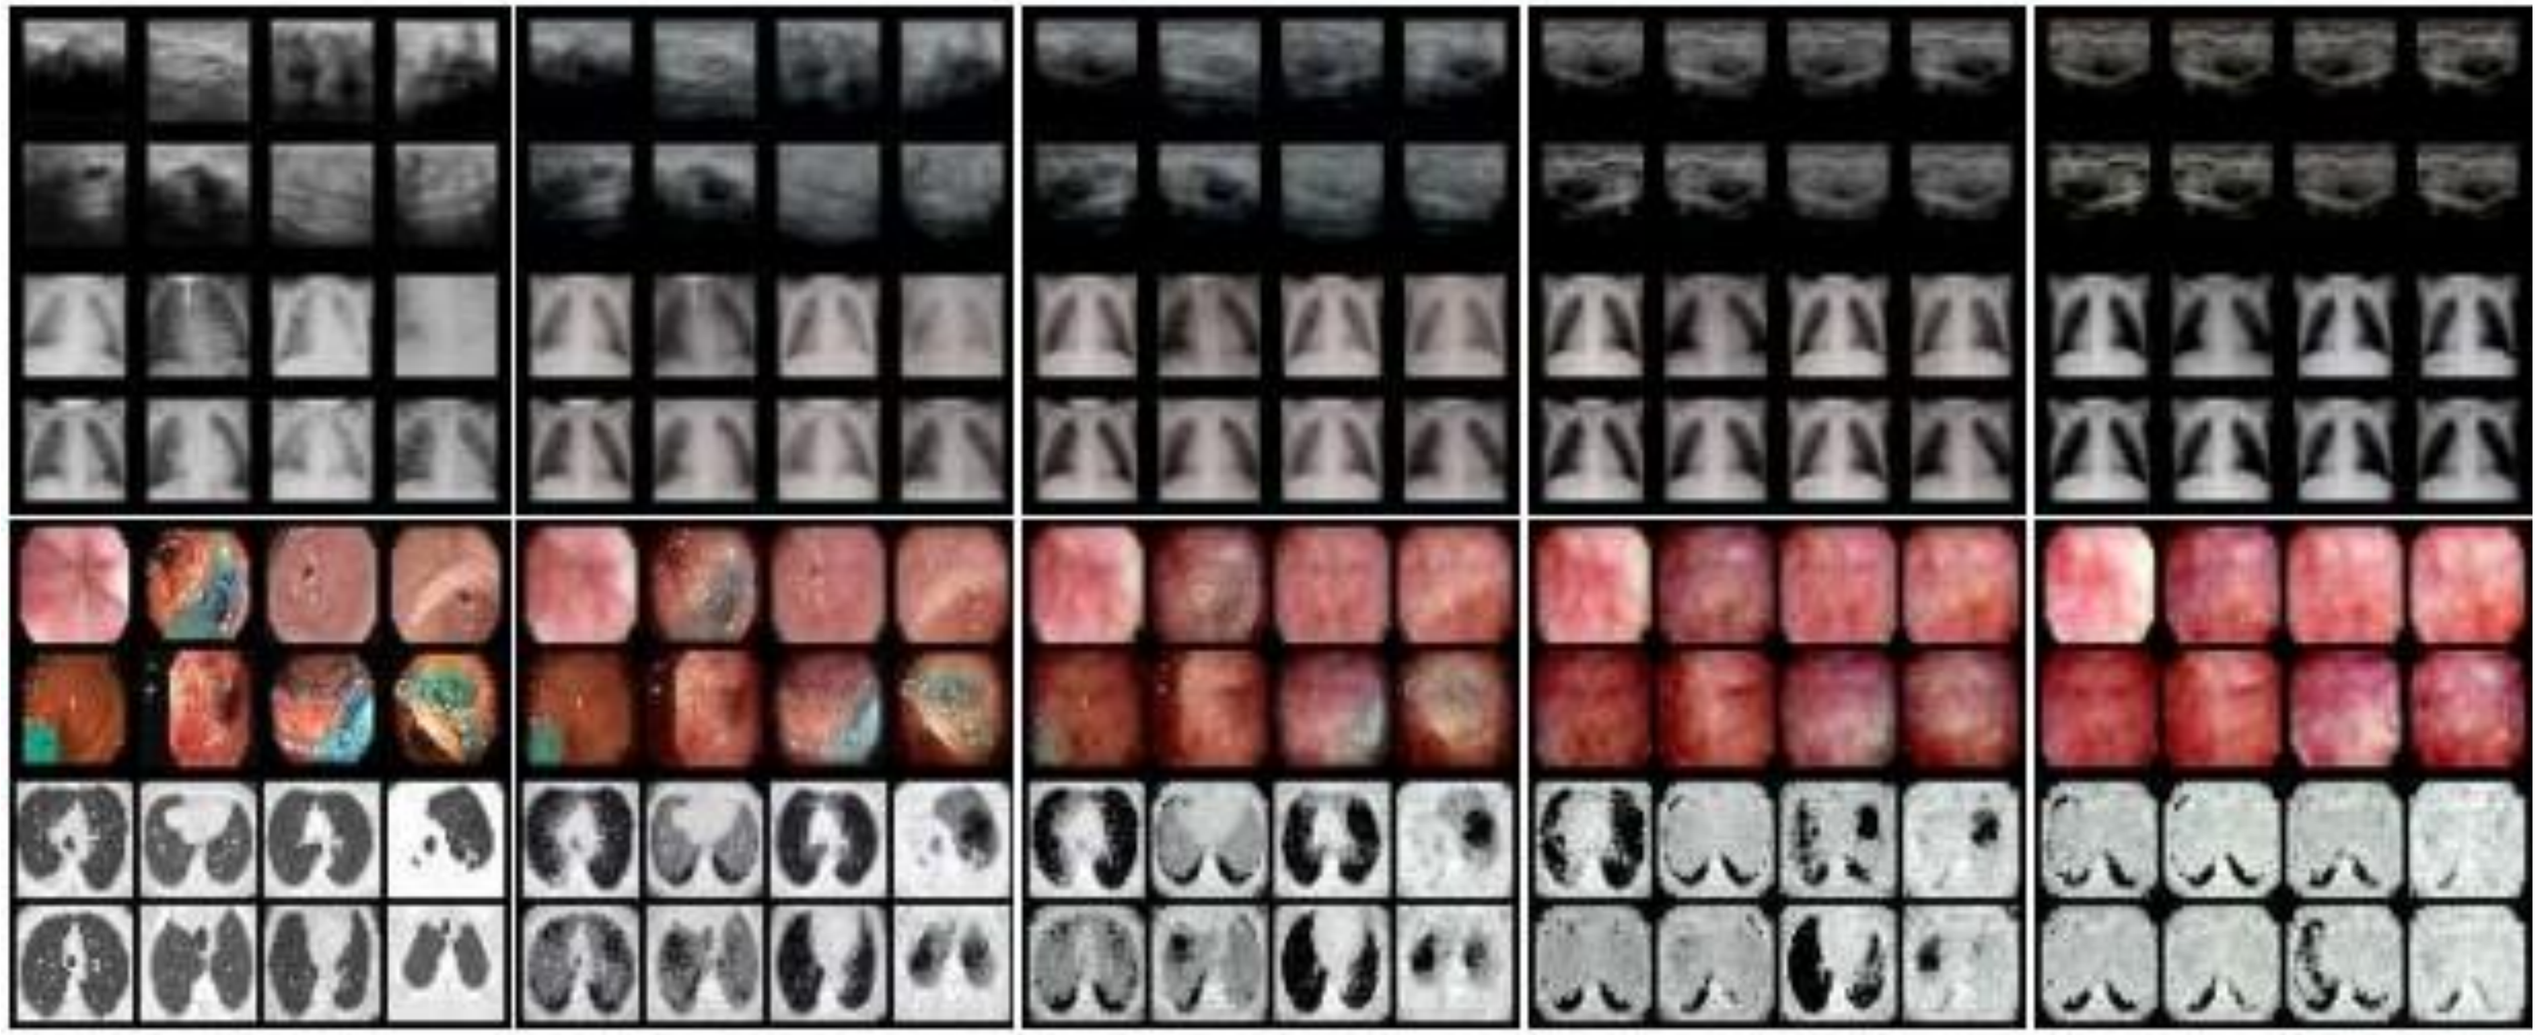

(a) Input Images (b) $t = 50$ (c) $t = 100$ (d) $t = 200$ (e) $t = 400$

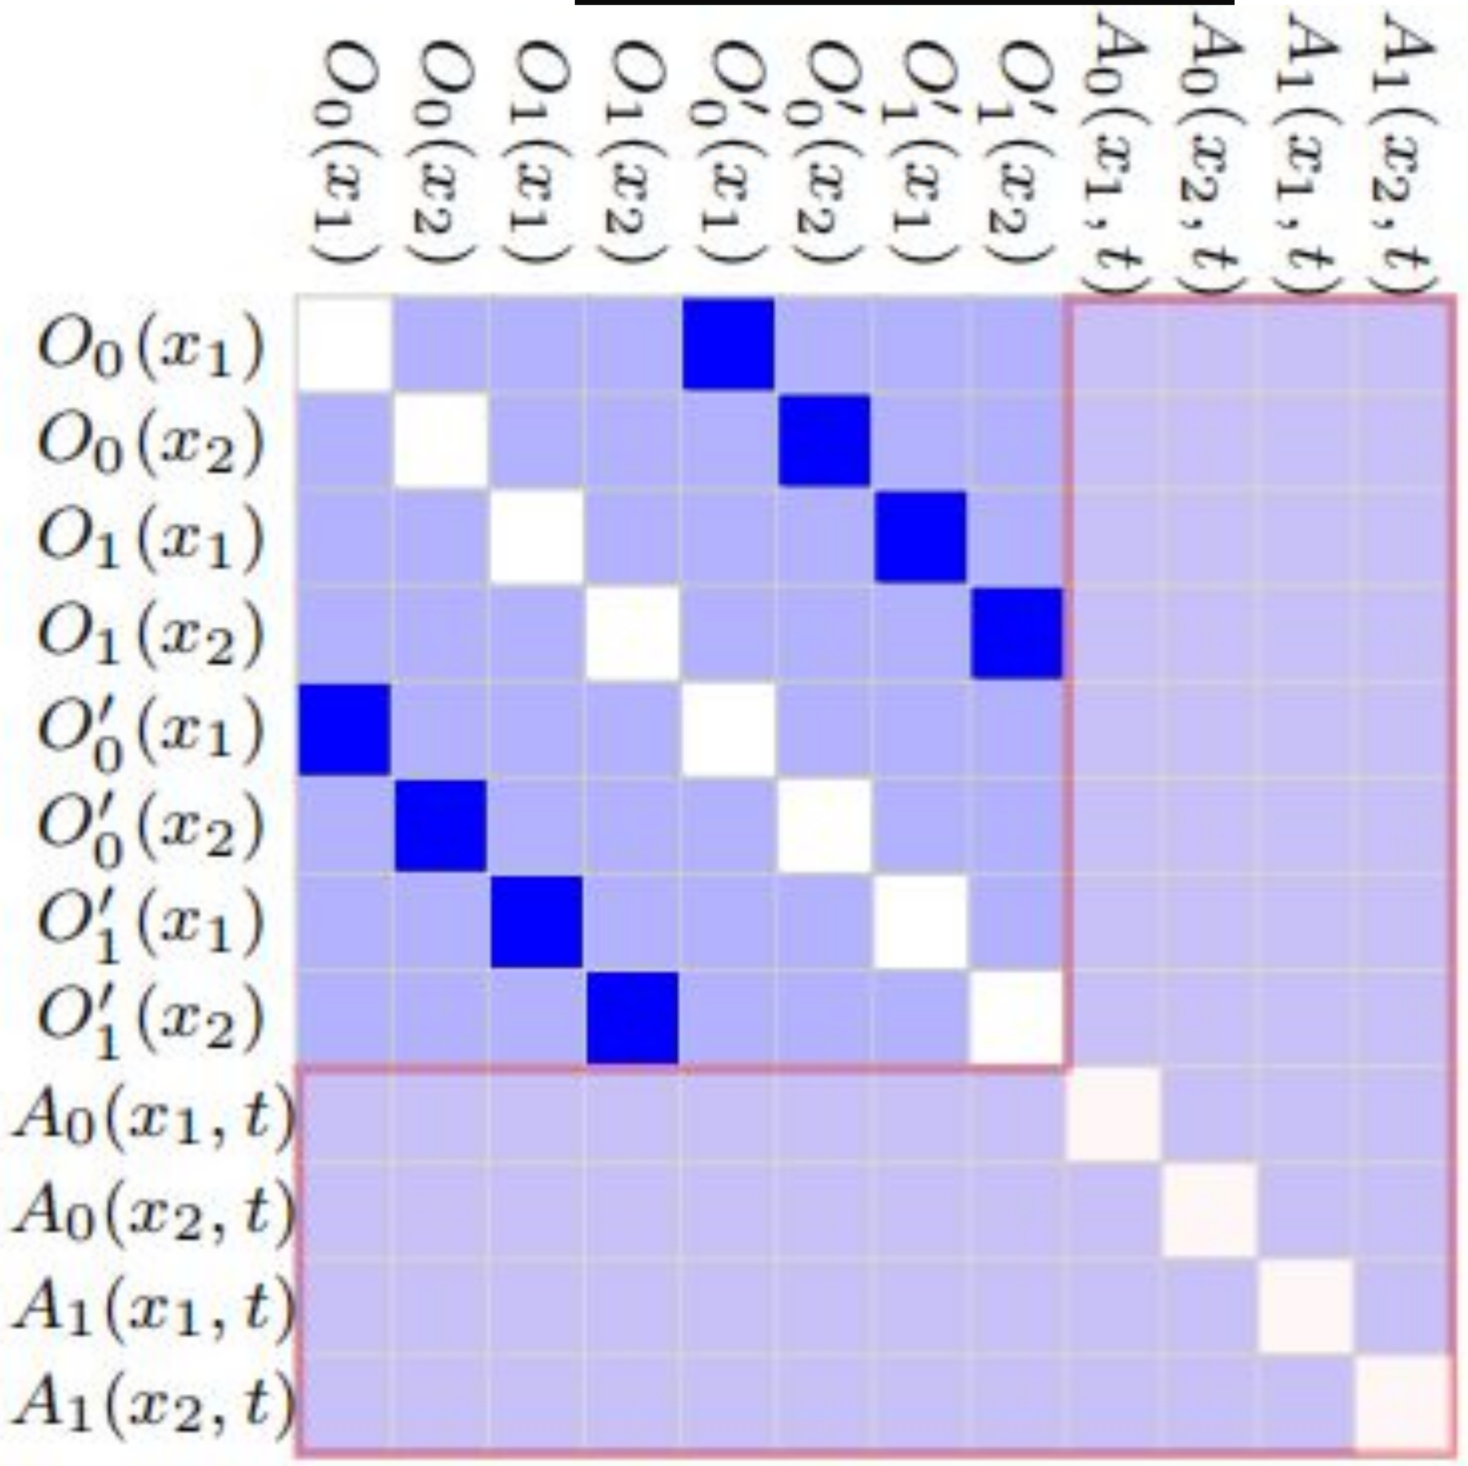

Key Technology:

Fine-Grained Contrastive Learning

Visualization of the target similarity matrix (K = 2 with two samples in a batch). The white, blue, and lavender blocks denote the excluded, positive, and negative pairs, respectively.

The red area contains the newly introduced negative pairs with dissolving transformations.

The proposed contrastive loss:

$$\mathcal{L}_{con} = \frac{1}{3BK}\frac{1}{|\{x_+\}|}\sum \ell_{i,j} \cdot \begin{cases} 0 & \mathbb{1}_{i,j} \in \{x_-\} \\ 1 & \mathbb{1}_{i,j} \in \{x_+\} \end{cases},$$

The DIA Framework

The DIA framework is designed to improve fine-grained anomaly detection in medical imaging, providing a robust solution to identify irregularities.

Our framework amplifies fine-details by dissolving the fine details.

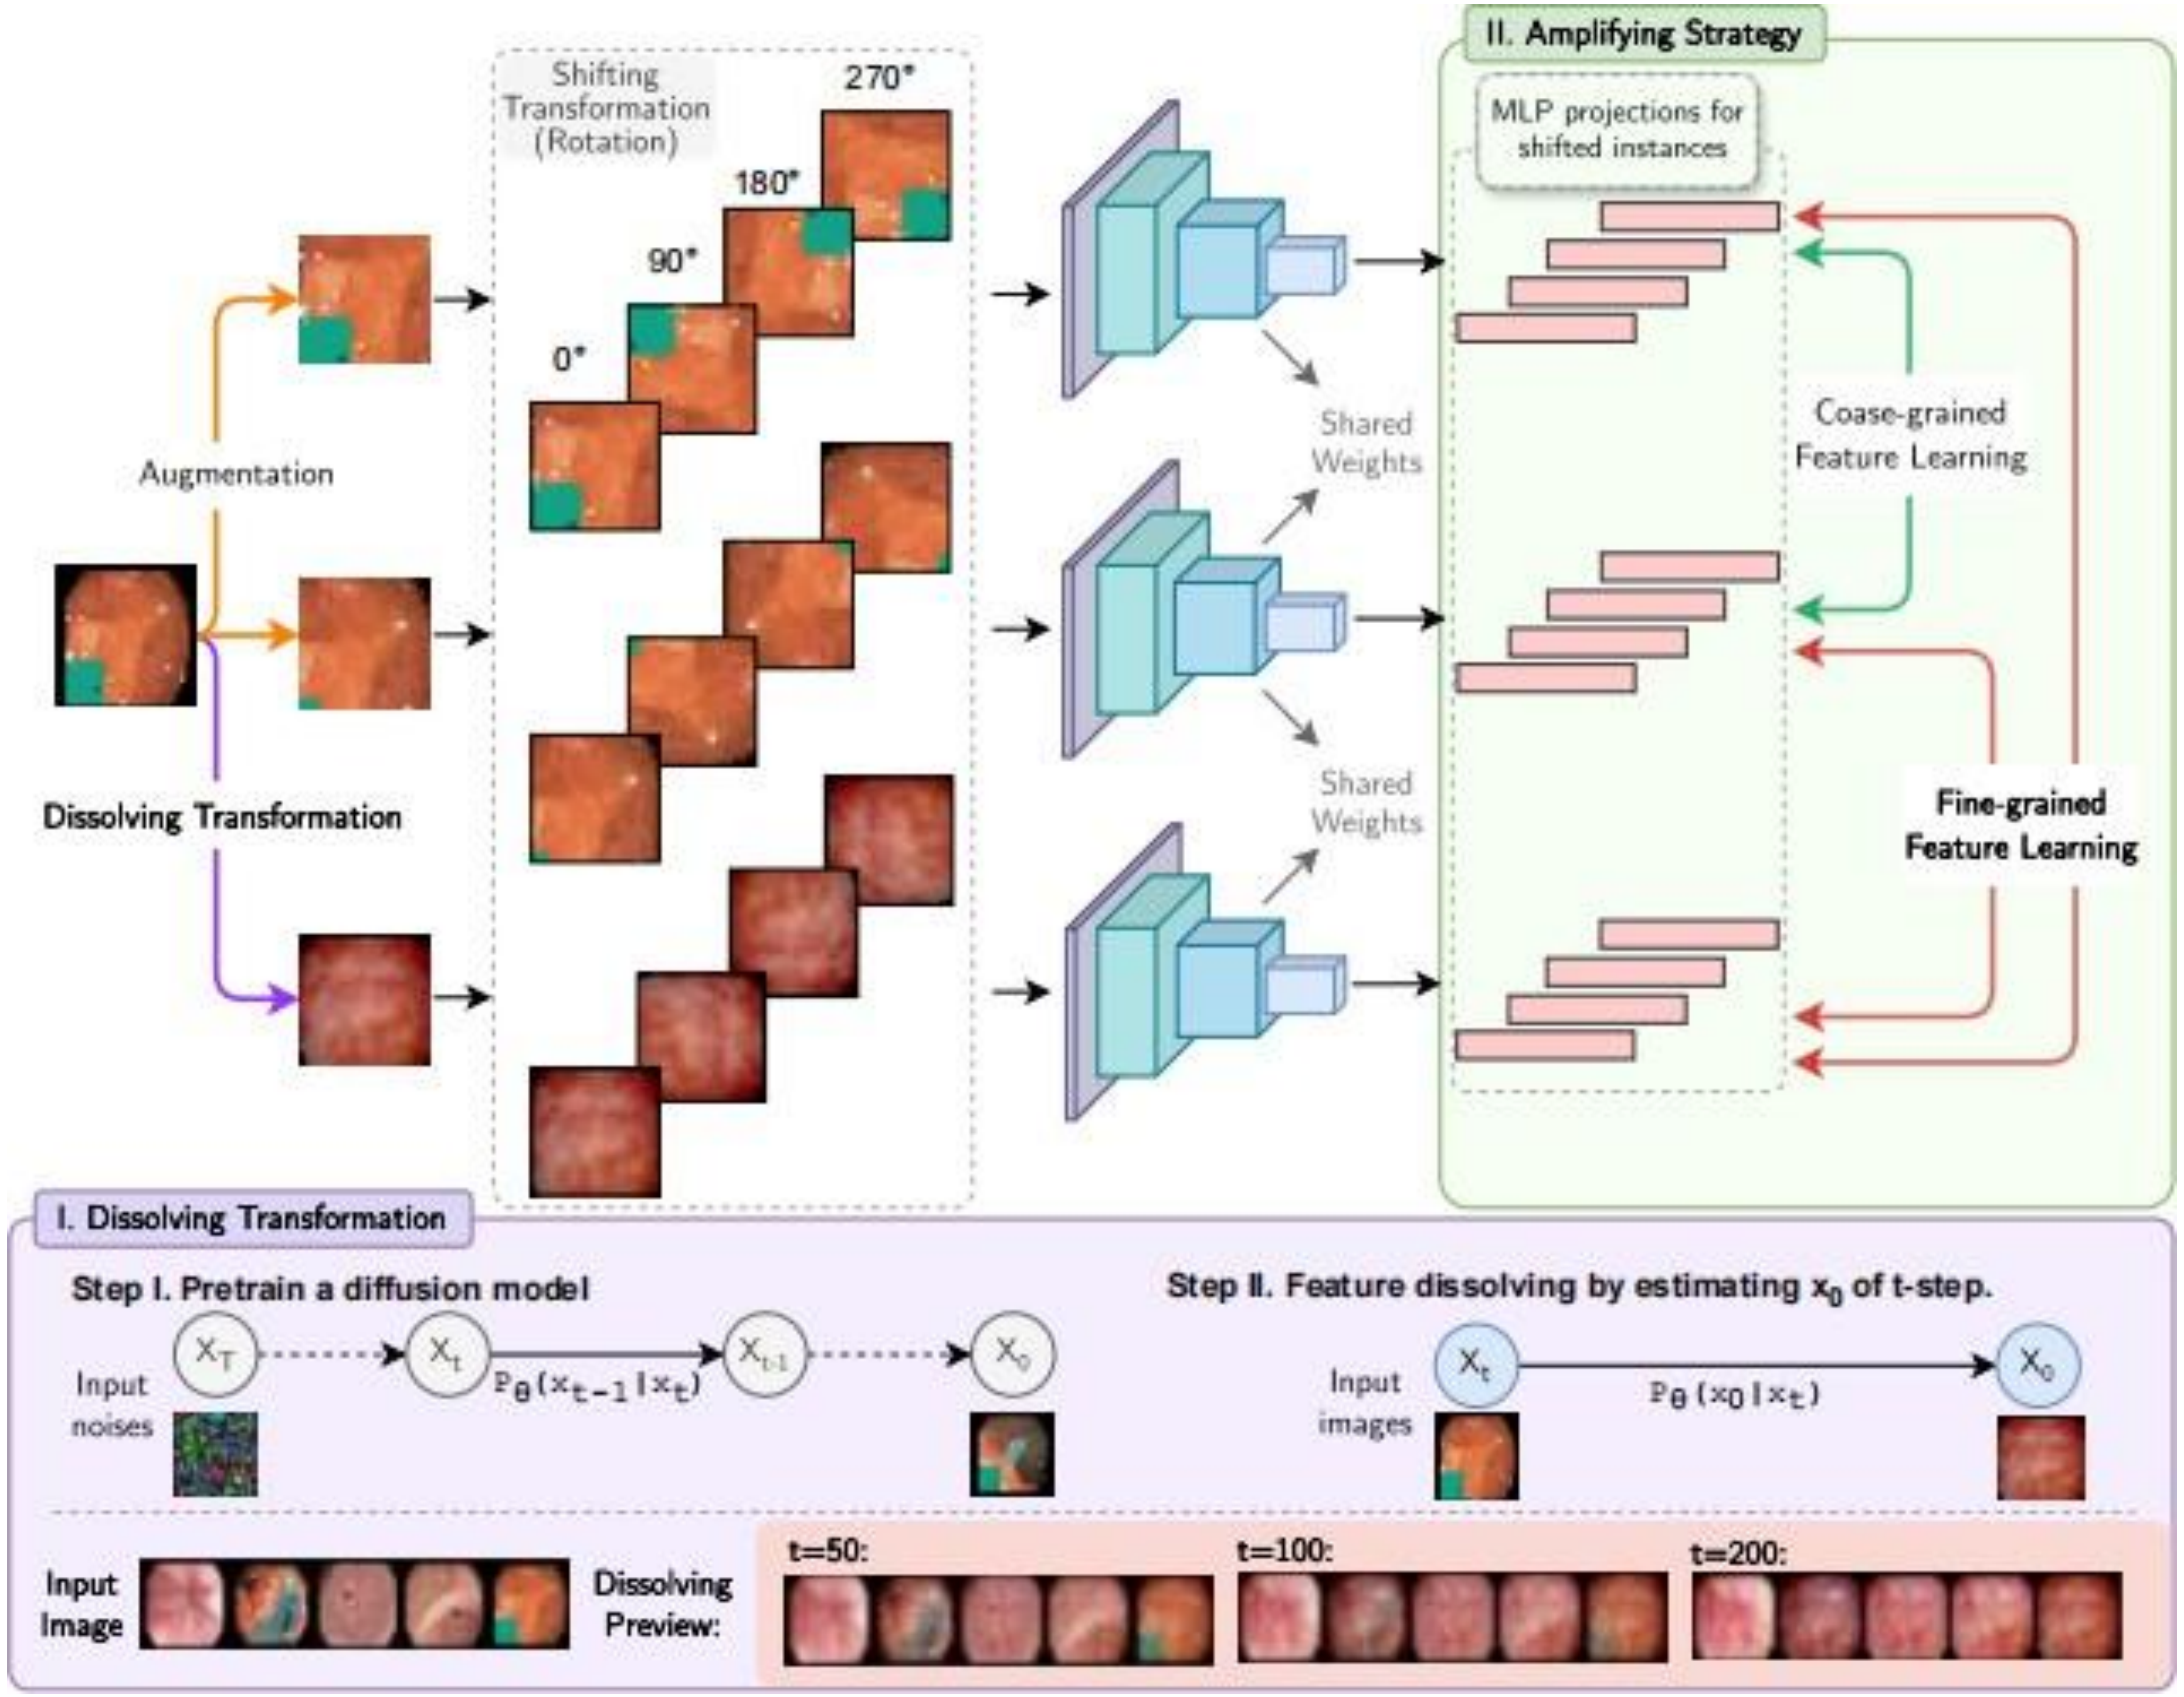

Significant AUC Improvement

Demonstrated a substantial AUC improvement of 18.40% over baseline methods across six diverse medical datasets.

Methods		Extra Training Data	Pnuemonia MNIST	Breast MNIST	SARS-COV-2	Kvasir Polyp	Retinal OCT	APTOS 2019
Reconstruction-based Methods								
GANomaly	(ACCV 18)	×	0.552$\pm$0.01	0.527$\pm$0.01	0.604$\pm$0.00	0.604$\pm$0.00	0.505$\pm$0.00	0.601$\pm$0.01
‡UniAD [60]	(NeurIPS 22)	×	0.734$\pm$0.02	0.624$\pm$0.01	0.636$\pm$0.00	0.724$\pm$0.03	0.921$\pm$0.01	0.874$\pm$0.00
Normalizing flow-based Methods								
‡CFlow [25]	(WACV 22)	×	0.537$\pm$0.01	0.647$\pm$0.01	0.622$\pm$0.01	0.852$\pm$0.03	0.712$\pm$0.02	0.452$\pm$0.01
UFlow [52]		×	0.792$\pm$0.01	0.631$\pm$0.01	0.653$\pm$0.02	0.562$\pm$0.02	0.630$\pm$0.01	0.731$\pm$0.00
FastFlow [61]		×	0.827$\pm$0.02	0.667$\pm$0.01	0.700$\pm$0.01	0.516$\pm$0.03	0.744$\pm$0.01	0.772$\pm$0.02
Teacher-Student Methods								
KDAD [45]	(CVPR 21)	×	0.378$\pm$0.02	0.611$\pm$0.02	0.770$\pm$0.01	0.775$\pm$0.01	0.801$\pm$0.00	0.631$\pm$0.01
RD4AD [18]	(CVPR 22)	✓	0.815$\pm$0.01	**0.759**$\pm$0.02	0.842$\pm$0.00	0.757$\pm$0.01	**0.996**$\pm$0.00	0.921$\pm$0.00
†Transformly [17]	(CVPR 22)	✓	0.821$\pm$0.01	0.738$\pm$0.04	0.711$\pm$0.00	0.568$\pm$0.00	0.824$\pm$0.01	0.616$\pm$0.01
‡EfficientAD [5]	(CVPR 24)	✓	0.686$\pm$0.02	0.696$\pm$0.03	0.711$\pm$0.02	0.753$\pm$0.03	0.826$\pm$0.02	0.763$\pm$0.02
Memory Bank-Based Methods								
CFA	(IEEE Access 22)	×	0.716$\pm$0.01	0.678$\pm$0.02	0.424$\pm$0.03	0.354$\pm$0.01	0.472$\pm$0.01	0.796$\pm$0.01
PatchCore	(CVPR 22)	×	0.737$\pm$0.01	0.700$\pm$0.02	0.654$\pm$0.01	0.832$\pm$0.01	0.758$\pm$0.01	0.583$\pm$0.01
Contrastive Learning-Based Methods								
Meanshift [41]	(AAAI 23)	×	0.818$\pm$0.02	0.648$\pm$0.01	0.767$\pm$0.03	0.694$\pm$0.05	0.438$\pm$0.01	0.826$\pm$0.01
CSI [51] Baseline	(NeurIPS 20)	×	0.834$\pm$0.03	0.546$\pm$0.03	0.785$\pm$0.02	0.609$\pm$0.03	0.803$\pm$0.00	0.927$\pm$0.00
DIA Ours		×	**0.903**$\pm$0.01	0.750$\pm$0.03	**0.851**$\pm$0.03	**0.860**$\pm$0.04	0.944$\pm$0.00	**0.934**$\pm$0.00

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

NEC
NEC Laboratories China

Thanks for Listening!