Toward Long-Tailed Online Anomaly Detection through Class-Agnostic Concepts

Chiao-An Yang[1] Kuan-Chuan Peng[2] Raymond A. Yeh[1]

1 Purdue University, Department of Computer Science

2 Mitsubishi Electric Research Laboratories

ICCV Honolulu Hawaii, Oct 19-23, 2025

TLDR: A new benchmark for ***Long-tailed online anomaly detection***. A ***class-agnostic*** AD framework that requires no class information.

Introduction

Anomaly detection (AD)

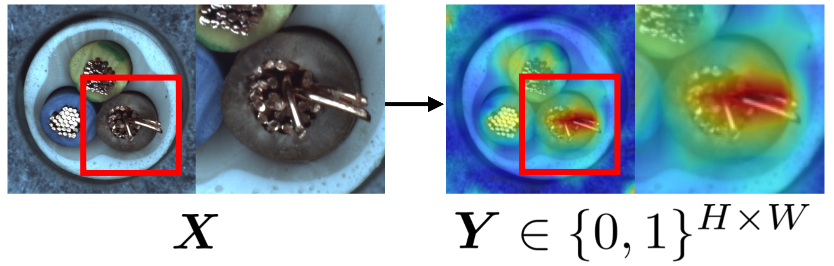

Long-tailed AD (LTAD)

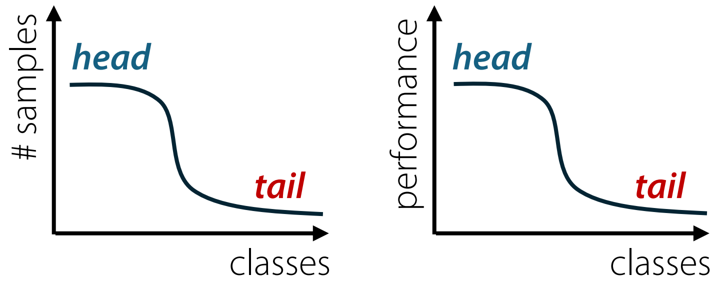

Class-aware AD with VLM

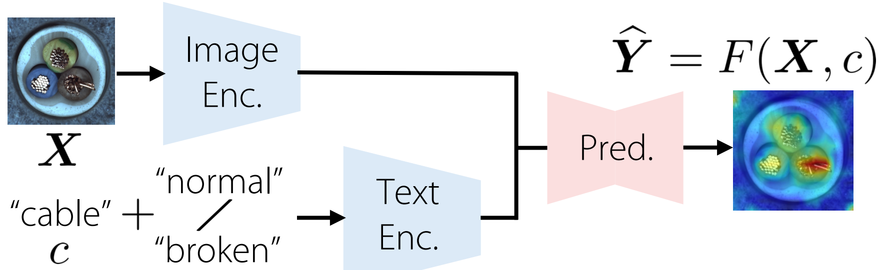

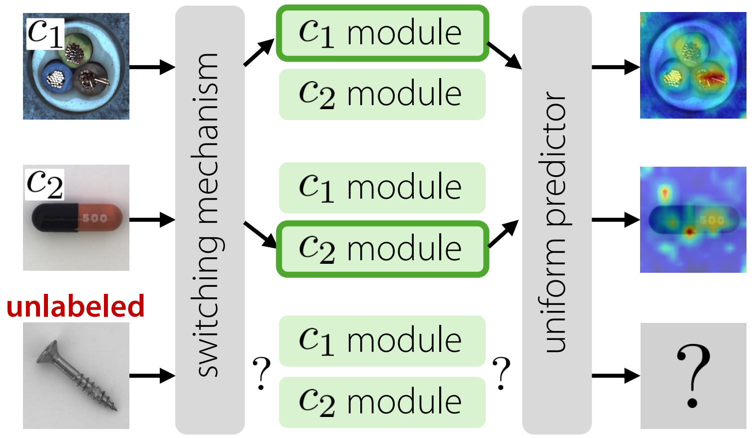

Approach (1)

Class-agnostic AD with concepts

We solve the problem of class-aware AD by learning a concept set $\widehat{\mathcal{C}}$ to approximate $\mathcal{C}$. When given images with an unseen class, our LTOAD can weight input images with concept-specific modules, i.e. $\{p_{\hat{c}}\}_{\hat{c}\in\widehat{\mathcal{C}}}$.

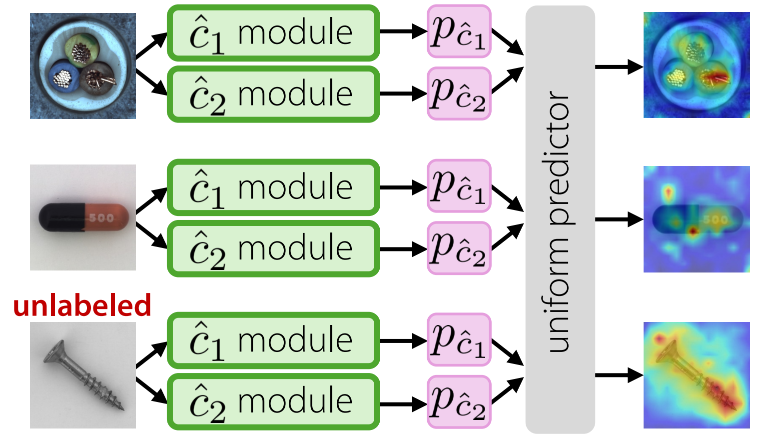

Making LTAD class-agnostic

The concept set $\widehat{\mathcal{C}}$ should be representative enough to cover the image classes $\mathcal{C}$ of interest. We learn $\widehat{\mathcal{C}}$ by measuring the pairwise similarity of images and vocabularies, plus majority voting.

At test time, a soft label is predicted for each image,

$$\boldsymbol{p} = \texttt{Softmax}(\{\langle \boldsymbol{f}, \boldsymbol{t}_{\hat{c}}\rangle\}_{\hat{c}\in\mathcal{C}})$$

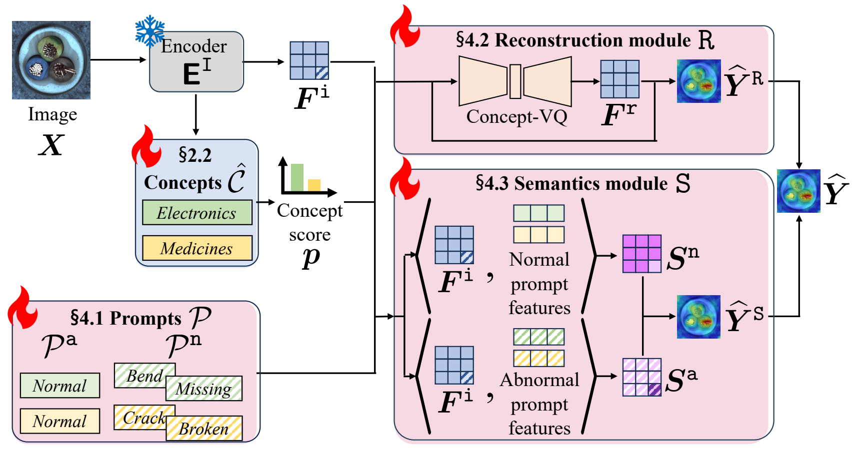

Approach (2)

Long-tailed online AD (LTOAD)

We propose an Anomalous Adaptive learning algorithm $\mathcal{A}^{\mathtt{AA}}$ to tackle the issues of overfitting and forgetting.

Output: $\Theta = [\theta_1, \cdots, \theta_T]$, where $T = |\mathcal{D}^{\mathtt{O}}|/\Delta$

Initialize $\Theta = [\,]$, $\widetilde{\theta}_0 = \theta_0$
for $t = 1$ to T **do**
  Acquire batch $\mathcal{D}^{\mathtt{O}}_t = \left[\widetilde{\boldsymbol{X}}_i\right]_{i=(t-1)\Delta,\cdots,t\Delta}$
  for $j = 1$ to Δ **do**
    Forward pass $\widehat{\boldsymbol{Y}}_j = \mathtt{F}_{\widetilde{\theta}_{t-1}}(\mathcal{D}^{\mathtt{O}}_t[j])$
    Compute pseudo abnormal label $\widehat{\boldsymbol{M}}_j = \mathcal{T}(\widehat{\boldsymbol{Y}}_j)$
    Compute mask-dependent loss $\widetilde{\mathcal{L}}_j = \widetilde{\mathcal{L}}(\widehat{\boldsymbol{M}}_j)$
    Compute gradient weight $\lambda_j \leftarrow \beta$ if $r(\widehat{\boldsymbol{Y}}_j) \geq \tau$ else 1
  end for
  $\widetilde{\theta}_t \leftarrow \widetilde{\theta}_{t-1} + \frac{1}{\Delta}\sum_j \lambda_j \nabla\widetilde{\mathcal{L}}_j$
  Update with EMA $\theta_t \leftarrow \gamma\theta_{t-1} + (1-\gamma)\widetilde{\theta}_t$
  Θ.append(θ_t)
end for
return Θ

Benchmarking LTOAD

The benchmark combines different session types ∈ {***blurry***, ***disjoint***} and different ordering types ∈ {***head-first***, ***tail-first***, ***mixed***}

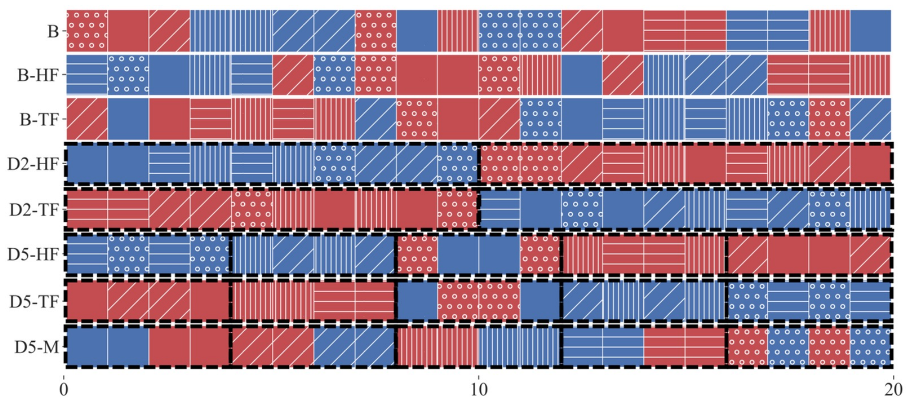

Experiments

Long-tailed offline AD (MVTec)

CA: class-agnostic

Method	CA	exp100	exp200	step100	step200
LTAD	x	94.46	94.18	93.83	92.12
HVQ	x	**95.25**	<u>94.79</u>	<u>94.17</u>	92.61
MoEAD	v	94.34	93.96	93.76	<u>92.76</u>
LTOAD	v	<u>95.21</u>	**94.94**	**95.11**	**94.00**

Concept set

Set ID	Domain	Size	Elements
$\mathcal{C}$	N/A	15	zipper, pill, capsule, grid, transistor, carpet, metal nut, wood, leather, screw, tile, cable, toothbrush, hazelnut, bottle*
LTOAD $\widehat{\mathcal{C}}$	In	10	semiconductor, zipper, beech, microscopy*, antibiotics, medicines, circuit , mahogany, hardwood, walnut

Long-tailed Offline AD (MVTec)

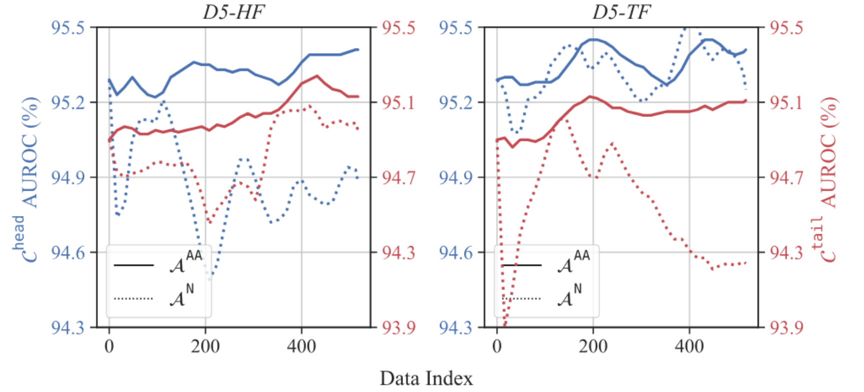

Cross-dataset evaluation

M: MVTec; ***V***: VisA; ***D***: DAGM. ***M > V***: Trained on M, tested on V.

Method	M > V	M > D	V > M	V > D	D > M	D > V
LTOAD	84.38	91.27	82.73	79.16	79.51	81.13
LTOAD + $\mathcal{A}^{\mathrm{AA}}$	**92.37**	**96.87**	**89.84**	**96.79**	**88.37**	**92.41**