

# Unsupervised Continual Anomaly Detection with Contrastively-Learned Prompt

Jiaqi Liu<sup>1\*</sup>, Kai Wu<sup>2\*</sup>, Qiang Nie<sup>2</sup>, Ying Chen<sup>2</sup>, Bin-Bin Gao<sup>2</sup>,  
Yong Liu<sup>2</sup>, Jinbao Wang<sup>1</sup>, Chengjie Wang<sup>2,3</sup>, Feng Zheng<sup>1†</sup>

<sup>1</sup>Southern University of Science and Technology

<sup>2</sup>Tencent YouTu Lab

<sup>3</sup>Shanghai Jiao Tong University

liujq32021@mail.sustech.edu.cn, lloydwu@tencent.com, stephennie@tencent.com, mumuychen@tencent.com,  
csgaobb@gmail.com, choasliu@tencent.com, linkingring@163.com, jasoncjwang@tencent.com, zfeng02@gmail.com

## Abstract

Unsupervised Anomaly Detection (UAD) with incremental training is crucial in industrial manufacturing, as unpredictable defects make obtaining sufficient labeled data infeasible. However, continual learning methods primarily rely on supervised annotations, while the application in UAD is limited due to the absence of supervision. Current UAD methods train separate models for different classes sequentially, leading to catastrophic forgetting and a heavy computational burden. To address this issue, we introduce a novel Unsupervised Continual Anomaly Detection framework called **UCAD**, which equips the UAD with continual learning capability through contrastively-learned prompts. In the proposed UCAD, we design a Continual Prompting Module (CPM) by utilizing a concise key-prompt-knowledge memory bank to guide task-invariant ‘anomaly’ model predictions using task-specific ‘normal’ knowledge. Moreover, Structure-based Contrastive Learning (SCL) is designed with the Segment Anything Model (SAM) to improve prompt learning and anomaly segmentation results. Specifically, by treating SAM’s masks as structure, we draw features within the same mask closer and push others apart for general feature representations. We conduct comprehensive experiments and set the benchmark on unsupervised continual anomaly detection and segmentation, demonstrating that our method is significantly better than anomaly detection methods, even with rehearsal training. The code will be available at <https://github.com/shirowalker/UCAD>.

## Introduction

Unsupervised Anomaly Detection (UAD) focuses on identifying unusual patterns or outliers in data without prior knowledge or labeled instances, relying solely on the inherent distribution of the ‘normal’ data (Chandola, Banerjee, and Kumar 2009). This approach is particularly useful in industrial manufacturing since acquiring well-labeled defect data can be challenging and costly.

Recent researches on UAD involve training distinct models for various classes, which inevitably relies on the knowledge of class identity during the test phase (Liu et al.

\*Contributed Equally.

†Corresponding Author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

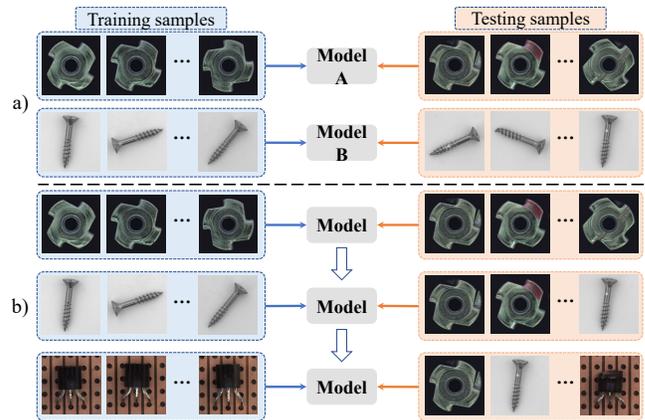


Figure 1: Comparison between separate models and UCAD methods: a) Using separate methods, each task has its own individual model. On the contrary, Ours b) uses a single model to handle all tasks without task identities. In the continuous stream, UCAD only requires the dataset of the current task for training and can be applied to previous tasks.

2023b). Moreover, forcing separate models to learn sequentially also results in a heavy computational burden with class incrementation. Some other methods focus on training a unified model that can handle multiple classes, such as UniAD (You et al. 2022). In real production, trains occur sequentially, which makes it impractical for UniAD to require all data to be trained simultaneously. Additionally, the unified model still lacks the ability to retain previously learned knowledge when continuously adapting to frequent product alterations during sequential training. Catastrophic forgetting and computational burden hinder UAD methods from applying to real-world scenarios.

Continual Learning (CL) is well-known for addressing the issue of catastrophic forgetting with a single model, especially when previous data is unavailable due to privacy reasons (Li et al. 2023). Recent research on continual learning can be categorized based on the requirement of task identities during the test phase. Task-aware approaches explicitly use the task identities to guide the learning process and prevent interference between tasks (Aljundi et al. 2018; Kirkpatrick et al. 2017). However, it’s not always possible to ac-

quire task identities during inference. Hence, task-agnostic methods are necessary and more prevail. Aljundi, Kelchtermans, and Tuytelaars progressively modifies data distribution to adapt various tasks in an online setup. L2P (Wang et al. 2022) dynamically learns prompts as task identities. Despite the effectiveness of task-agnostic CL methods in supervised tasks, their efficacy in UAD remains unproven. Obtaining large scales of anomalous data is difficult in industries due to high production success rates and privacy concerns. Therefore, it is crucial to explore the application of CL in UAD.

To date, there is no known effort, except for Gaussian distribution estimator (DNE) (Li et al. 2022), to incorporate CL into UAD. However, DNE still relies on augmentations (Li et al. 2021) to provide pseudo-supervision and is not applicable to anomaly segmentation. DNE can be considered a continual binary image classification method rather continual anomaly detection (AD) method. In real industrial manufacturing, accurately segmenting the areas of anomalies is essential for anomaly standard quantization. Hence, there is an urgent need for a method that can perform unsupervised continual AD and segmentation simultaneously.

To address the aforementioned problems, we propose a novel framework for Unsupervised Continual Anomaly Detection called UCAD, which can sequentially learn to detect anomalies of different classes using a single model, as shown in Fig. 1. UCAD incorporates a Continual Prompting Module (CPM) to enable CL in unsupervised AD and a Structure-based Contrastive Learning (SCL) module to extract more compact features across various tasks. The CPM learns a “key-prompt-knowledge” memory space to store auto-selected task queries, task adaptation prompts, and the ‘normal’ knowledge of different classes. Given an image, the key is automatically selected to retrieve the corresponding task prompts. Based on the prompts, the image feature is further extracted and compared with its normal knowledge for anomaly detection, similar to PatchCore (Roth et al. 2022). However, the performance of CPM is limited because the frozen backbone (ViT) cannot provide compact feature representations across various tasks. To overcome this limitation, the SCL is introduced to extract more dominant feature representations and reduce domain gaps by leveraging the general segmentation ability of SAM (Kirillov et al. 2023). With SCL, features of the same structure (segmented area) are pulled together and pushed away from features in other structures. As a result, the prompts are contrastively learned for better feature extraction across different tasks. Our contributions can be summarized as follows:

- To the best of our knowledge, our proposed UCAD is the first framework for task-agnostic continual learning on unsupervised anomaly detection and segmentation. UCAD novelty learns a key-prompt-knowledge memory space for automatic task instruction, knowledge transfer, unsupervised anomaly detection and segmentation.
- We propose to use contrastively-learned prompts to improve unsupervised feature extraction among various classes by exploiting the general capabilities of SAM.
- We have conducted thorough experiments and intro-

duced a new benchmark for unsupervised CL anomaly detection and segmentation. Our proposed UCAD outperforms previous state-of-the-art (SOTA) AD methods by 15.6% on detection and 26.6% on segmentation.

## Related Work

### Unsupervised Image Anomaly Detection

With the release of the MVTEC AD dataset (Bergmann et al. 2019), the development of industrial image anomaly detection has shifted from a supervised paradigm to an unsupervised paradigm. In the unsupervised anomaly detection paradigm, the training set only consists of normal images, while the test set contains both normal images and annotated abnormal images. Gradually, research on unsupervised industrial image anomaly detection has been divided into two main categories: feature-embedding-based methods and reconstruction-based methods (Liu et al. 2023b). **Feature-embedding-based methods** can be further categorized into four subcategories, including *teacher-student model* (Bergmann et al. 2020; Salehi et al. 2021; Deng and Li 2022; Tien et al. 2023), *one-class classification methods* (Li et al. 2021; Liu et al. 2023c), *mapping-based methods* (Rudolph, Wandt, and Rosenhahn 2021; Gudovskiy, Ishizaka, and Kozuka 2022; Rudolph et al. 2022; Lei et al. 2023) and *memory-based methods* (Defard et al. 2021; Roth et al. 2022; Jiang et al. 2022b; Xie et al. 2022; Liu et al. 2023a). **Reconstruction-based methods** can be categorized based on the type of reconstruction network, including *autoencoder-based methods* (Zavrtanik, Kristan, and Skočaj 2021, 2022; Schlüter et al. 2022), *Generative Adversarial Networks (GANs)* (Goodfellow et al. 2014) *based methods* (Yan et al. 2021; Liang et al. 2022), *ViT-based methods* (Mishra et al. 2021; Pirnay and Chai 2022; Jiang et al. 2022a), and *Diffusion model-based methods* (Mousakhan, Brox, and Tayyub 2023; Zhang et al. 2023).

However, existing UAD methods are designed to enhance AD capabilities within a single category. They often lack the ability to perform anomaly detection in a continual learning scenario. Our method is specifically designed for the scenario of continual learning and achieves continual anomaly segmentation in an unsupervised manner.

### Continual Image Anomaly Detection

Different from natural image object detection tasks, the data stream is common in industrial manufacturing. Some current methods have recognized this phenomenon and attempted to design algorithms specifically to address the challenges in this scenario. IDDM (Zhang and Chen 2023) presents an incremental anomaly detection method based on a small number of labeled samples. On the other hand, LeMO (Gao et al. 2023) follows the common unsupervised anomaly detection paradigm and performs incremental anomaly detection as normal samples continuously increase. However, both IDDM and LeMO focus on intra-class continual anomaly detection research without addressing inter-class incremental anomaly detection challenges. Li et al.’s research (Li et al. 2022) is the most closely related to ours. They propose DNE

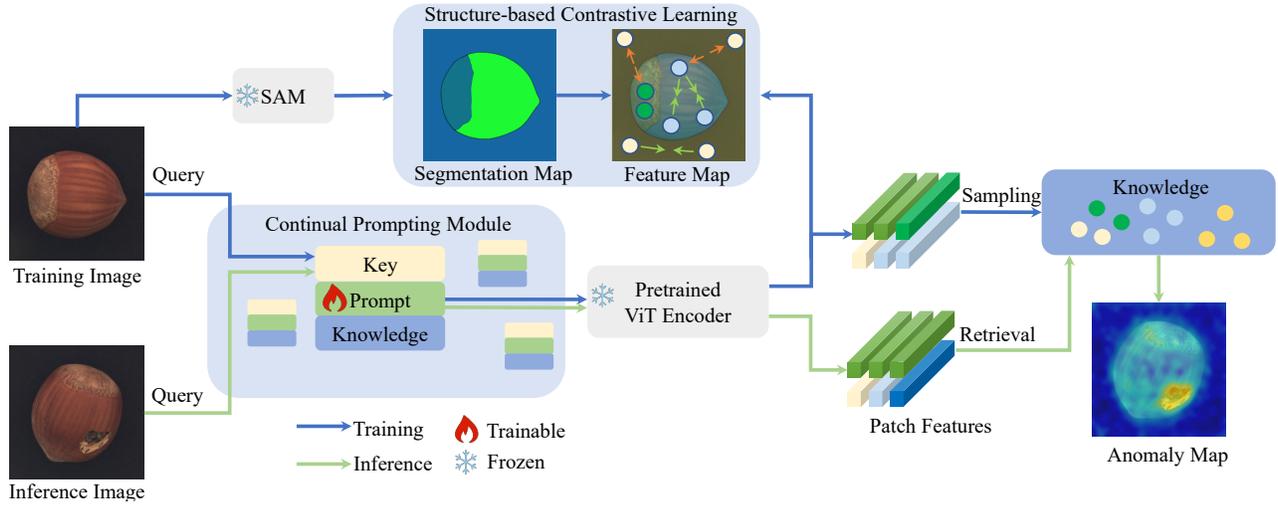


Figure 2: The framework of UCAD mainly comprises a Continual Prompting Module (CPM) and a Structure-based Contrastive Learning (SCL) module, integrated with the SAM network. During training, the CPM establishes a key-prompt-knowledge system that efficiently maintains training data information, while also reducing memory and computational resource usage. Moreover, UCAD proposes a contrastive learning method using the SAM segmentation map to enhance the feature representations. Finally, the detection of anomalies is accomplished by comparing current features and retrieved task-specific knowledge.

for image-level anomaly detection in continual learning scenarios. Due to the limitation of DNE in storing only class-level information, it cannot perform fine-grained localization, thus making it unsuitable for anomaly segmentation. Our method goes beyond continual anomaly classification and extends to pixel-level continual anomaly detection.

## Methods

### Unsupervised Continual AD Problem Definition

Unsupervised Anomaly Detection (AD) aims to identify anomalous data using only normal data, since obtaining labeled anomalous samples is challenging in real-world industrial production scenarios. The training set contains only normal samples from various tasks, while the test set includes both normal and abnormal samples, reflecting real-world applications. To formulate the problem, we define the multi-class training set as  $\mathcal{T}_{train}^{total} = \{\mathcal{T}_{train}^1, \mathcal{T}_{train}^2, \dots, \mathcal{T}_{train}^n\}$  and test set as  $\mathcal{T}_{test}^{total} = \{\mathcal{T}_{test}^1, \mathcal{T}_{test}^2, \dots, \mathcal{T}_{test}^n\}$ .  $\mathcal{T}_{train}^i$  and  $\mathcal{T}_{test}^i$  represent class  $i$ -th training and test data, respectively.

Under unsupervised continual AD and segmentation setting, a unified model is trained non-repetitively on incrementally added classes. Given  $N_{task}$  tasks or classes, the model is sequentially trained on sub-training sets  $\mathcal{T}_{train}^i, i \in N_{task}$  and subsequently tested on all past test subdatasets  $\mathcal{T}_{test}^{total}$ . This evaluation method ensures the final trained model's ability to retain previously acquired knowledge.

### Continual Prompting Module

Applying CL to unsupervised AD faces two challenges: 1) How to determine the task identities of the incoming image automatically; 2) How to guide the model's predictions for

the relevant task in an unsupervised manner. Thus, a continual prompting module is designed to dynamically adapt and instruct unsupervised model predictions. We propose to use a memory space  $\mathcal{M}$  for a **key-prompt-knowledge** architecture,  $(\mathcal{K}_e, \mathcal{V}, \mathcal{K}_n)$ , that contains two distinct phases: the task identification phase and the task adaptation phase.

In the task identification phase, images  $x \in \mathbb{R}^{H \times W \times C}$  will go through a frozen pretrained vision transformer  $f$  (ViT) to extract keys  $k \in \mathcal{K}_e$ , also known as task identities. Because task identity contains both textual details and high-level information, we use a specific layer of ViT rather than the last embedding  $k = f^i(x), k \in \mathbb{R}^{N_p \times C}$ , in which  $k$  is the feature and  $N_p$  is the num of patches after  $i$ -th block (in this paper, we use  $i = 5$ ). However, assuming we have  $N_I$  training images for task  $t$ , all extracted embeddings would have dimension  $\mathcal{K}_e^t \in \mathbb{R}^{N_I \times N_p \times C}$ , which means a lot of memory space. To make task matching efficient during testing, we propose to use one image's feature space representing the whole task  $\mathbb{R}^{N_I \times N_p \times C} \rightarrow \mathbb{R}^{N_p \times C}$ . Note that a single image's feature space is negligible compared to the whole task in the continual training setting. We find that the farthest point sampling method (Eldar et al. 1997) is efficient for selecting representative features to serve as **keys**. So task identities  $\mathcal{K}_e$  can be represented as a set:

$$\begin{aligned} \mathcal{K}_e^t &= FPS(\mathcal{K}_e^t), \quad \mathcal{K}_e^t \in \mathbb{R}^{N_p \times C} \\ \mathcal{K}_e &= \{\mathcal{K}_e^0, \mathcal{K}_e^1, \dots, \mathcal{K}_e^t\}, \quad t \in N_{task}, \end{aligned} \quad (1)$$

where  $FPS$  is furthest point sampling,  $\mathcal{K}_e^t$  represents all extracted embeddings of task  $t$ .

During the task adaptation phase, inspired by (Liu et al. 2021) which injects new knowledge into models, we design learnable **prompts**  $\mathcal{V}$  to transfer task-related information to the current image. Unlike  $\mathcal{K}_e$  that is downsampled from the pretrained backbone, prompts  $p \in \mathcal{V}$  are purely learnable

to accommodate the current task. We add a prompt  $p^i$  to each layer’s input feature to convey task information to the current image,  $k^i = f^i(k^{i-1} + p^i)$ , where  $k^i$  is the output feature of the  $i$ -th layer,  $k^{i-1}$  is the input feature, and  $p^i$  is the prompt added to the  $i$ -th layer to transfer task-specific information to the current image. Then, the task-transferred image features  $k^i$  are used to create the **knowledge**  $\mathcal{K}_n$  during training. Since we are not using supervision,  $\mathcal{K}_n$  serves as the standard to distinguish anomaly data by comparing it to the test image features. However, image features can be exceedingly large during training accumulation, we use coreset sampling (Roth et al. 2022) to reduce storage for  $\mathcal{V}$ .

$$\begin{aligned} \mathcal{K}_n &= \text{CoreSetSampling}(k^i) \\ &= \arg \min_{\mathcal{M}_c \subset \mathcal{M}} \max_{m \in \mathcal{M}} \min_{n \in \mathcal{M}_c} \|m - n\|_2, \end{aligned} \quad (2)$$

where  $\mathcal{M}$  is the nominal image features during training,  $\mathcal{M}_c$  is the coreset space for patch-level features  $k^i$ , and  $i = 5$  in our experiments since middle features contain both context and semantic information. After establishing key-prompt-knowledge correspondance for each task, our proposed Continual Prompting Module can successfully transfer knowledge from previous tasks to the current image. However, the features stored in  $\mathcal{V}$  may not be discriminative enough because the backbone  $f$  has been pretrained and not adapted to the current task. To make feature representations more compact, we developed a structure-based contrastive learning method to learn prompt contrastively.

### Structure-Based Contrastive Learning

Inspired by ReConPatch (Hyun et al. 2023), we designed structure-based contrastive learning to enhance the network representation for patch-level comparison during testing. We discovered that SAM (Kirillov et al. 2023) consistently provides general structure knowledge, such as masks, without requiring training. As illustrated in Figure 2, for each image in the training set, we employed SAM to generate corresponding segmentation images  $I_s$ , in which different regions represent distinct structures or semantics. Simultaneously, guided by prompts, we obtain the feature map  $F_s \in k^i$  for each region, where  $k^i$  is the  $i$ -th layer feature in the previous section. We downsampled the segmentation image  $I_s$  to match the size of  $F_s \in \mathbb{R}^{c \times h \times w}$  and aligned the corresponding positions to create the label map  $L_s$ . By incorporating contrastive learning, the knowledge generality in  $\mathcal{K}_n$  is achieved by pulling the features of the same region closer and pushing the features of different regions further apart. The loss function is:

$$\begin{aligned} L_{pos.con} &= \sum_{i,p}^H \sum_{j,q}^W \cos(F_{ij}, F_{pq}), (L_{ij} = L_{pq}), \\ L_{neg.con} &= \sum_{i,p}^H \sum_{j,q}^W \cos(F_{ij}, F_{pq}), (L_{ij} \neq L_{pq}), \\ L_{total} &= \lambda_\alpha L_{neg.con} - \lambda_\beta L_{pos.con}. \end{aligned} \quad (3)$$

In the given paragraph,  $F_{ij}$  denotes the embedding of feature  $F_s$  at position  $(i, j)$  with a shape of  $(1, 1, c)$ , while  $L_{ij}$  repre-

sents the label of feature  $F_{ij}$  at the corresponding position in the segmentation result generated by SAM.  $\lambda_\alpha$  and  $\lambda_\beta$  are 1. By training prompts using this contrastive loss, the model’s representation ability is enhanced, and features of various textures become more compact. Consequently, this approach results in more distinct representations of abnormal features during testing, allowing them to stand out prominently.

### Test-Time Task-Agnostic Inference

**Task Selection and Adaption** To automatically determine the task identity during testing, an image  $x^{test}$  initially locates its corresponding task by selecting from  $\mathcal{K}_e$  based on the highest similarity. The corresponding task identity is selected by the equation below:

$$\begin{aligned} \mathcal{K}_e^t &= \arg \min_{m \in \mathcal{K}_e} \text{Sim}(m - m^{test}), \\ \text{Sim}(m - m^{test}) &= \sum_{x \in N_p} \min_{y \in N_p} \|m_x - m_y^{test}\|_2, \end{aligned} \quad (4)$$

where  $m^{test}$  is the patch-level feature from  $i$ -th layer feature map of ViT containing multiple patches  $N_p$ ,  $i = 5$  in this paper as discussed in previous section. Since the utilization of a key-prompt-knowledge architecture, the associated prompts  $\mathcal{V}$  and knowledge  $\mathcal{K}_n$  can be readily retrieved. By combining the selected prompts with test patches and processing them through ViT, features from the test sample are adapted and extracted. Subsequently, anomaly scores are calculated based on the minimum distance to the task’s knowledge  $\mathcal{K}_n^t$ .

**Anomaly Detection and Segmentation** To calculate the anomaly score, we compare the image feature  $m^{test}$  with the nominal features stored in task-specific knowledge base  $\mathcal{K}_n^t$ . Building upon the patch-level retrieval, we employed re-weighting to implement the anomaly detection process.  $\mathcal{N}_b(m^*)$  represents the nearest neighbors of  $m^*$  in  $\mathcal{K}_n^t$ . We use the distance between  $m^{test}$  and  $m^*$  as the basic anomaly score, and then calculate the distance between  $m^{test}$  and the features in  $\mathcal{N}_b(m^*)$  to achieve the re-weighting effect. Through Equation 5, we set the furthest distance between feature  $m^{test,*}$  in the test feature set  $\mathcal{P}(x^{test})$  and memory bank  $\mathcal{K}_n^t$  to represent the anomaly score  $s^*$  of the sample.

$$\begin{aligned} m^{test,*}, m^* &= \arg \max_{m^{test} \in \mathcal{P}(x^{test})} \arg \min_{m \in \mathcal{K}_n^t} \|m^{test} - m\|_2, \\ s^* &= \|m^{test,*} - m^*\|_2. \end{aligned} \quad (5)$$

By re-weighting from neighbors  $m^* \in \mathcal{K}_n^t$ , the anomaly score  $s$  becomes more robust, as in Equation 6:

$$s = \left( 1 - \frac{\exp \|m^{test,*} - m^*\|_2}{\sum_{m \in \mathcal{N}_b(m^*)} \exp \|m^{test,*} - m\|_2} \right) \cdot s^*. \quad (6)$$

The anomaly score of the image is calculated by the max score of all patches,  $S_{img} = \max(s_i), i \in N_p$ . The coarse segmentation map,  $S_{cmap}$ , is represented by scores calculated from each patch. By upsampling and applying Gaussian smoothing to  $S_{cmap}$ , the final segmentation result  $S_{map}$  is obtained with the same dimensions as the input image.

Methods	BO	CA	CAP	CAR	GR	HA	LE	MN	PI	SC	TI	TO	TR	WO	ZI
CFA	0.309	0.489	0.275	0.834	0.571	0.903	0.935	0.464	0.528	0.528	0.763	0.519	0.320	0.923	0.984
CSFlow	0.129	0.420	0.363	0.978	0.602	0.269	0.906	0.220	0.263	0.434	0.697	0.569	0.432	0.802	0.997
CutPaste	0.111	0.422	0.373	0.198	0.214	0.578	0.007	0.517	0.371	0.356	0.112	0.158	0.340	0.150	0.775
DRAEM	0.793	0.411	0.517	0.537	0.799	0.524	0.480	0.422	0.452	<b>1.000</b>	0.548	0.625	0.307	0.517	<b>0.996</b>
FastFlow	0.454	0.512	0.517	0.489	0.482	0.522	0.487	0.476	0.575	0.402	0.489	0.267	0.526	0.616	0.867
FAVAE	0.666	0.396	0.357	0.610	0.644	0.884	0.406	0.416	0.531	0.624	0.563	0.503	0.331	0.728	0.544
PaDiM	0.458	0.544	0.418	0.454	0.704	0.635	0.418	0.446	0.449	0.578	0.581	0.678	0.407	0.549	0.855
PatchCore	0.163	0.518	0.350	0.968	0.700	0.839	0.625	0.259	0.459	0.484	0.776	0.586	0.341	0.970	0.991
RD4AD	0.401	0.538	0.475	0.583	0.558	0.909	0.596	0.623	0.479	0.596	0.715	0.397	0.385	0.700	0.987
SPADE	0.302	0.444	0.525	0.529	0.460	0.410	0.577	0.592	0.484	0.514	0.881	0.386	0.622	0.897	0.949
STPM	0.329	0.539	0.610	0.462	0.569	0.540	0.740	0.456	0.523	0.753	0.736	0.375	0.450	0.779	0.783
SimpleNet	0.938	0.560	0.519	0.736	0.592	0.859	0.749	0.710	0.701	0.599	0.654	0.422	0.669	0.908	0.996
UniAD	0.801	0.660	0.823	0.754	0.713	0.904	0.715	0.791	0.869	0.731	0.687	0.776	0.490	0.903	0.997
DNE	0.990	0.619	0.609	0.984	<b>0.998</b>	0.924	<b>1.000</b>	<b>0.989</b>	0.671	0.588	0.980	0.933	0.877	0.930	0.958
PatchCore*	0.533	0.505	0.351	0.865	0.723	0.959	0.854	0.456	0.511	0.626	0.748	0.600	0.427	0.900	0.974
UniAD*	0.997	0.701	0.765	<b>0.998</b>	0.896	0.936	<b>1.000</b>	0.964	<b>0.895</b>	0.554	0.989	0.928	<b>0.966</b>	0.982	0.987
Ours	<b>1.000</b>	<b>0.751</b>	<b>0.866</b>	0.965	0.944	<b>0.994</b>	<b>1.000</b>	0.988	0.894	0.739	<b>0.998</b>	<b>1.000</b>	0.874	0.995	0.938

Table 1: Image-level AUROC $\uparrow$  on MVTec AD dataset (Bergmann et al. 2019) after training on the last subdataset. Note that \* signifies the usage of a cache pool for rehearsal during training which may not be possible in real applications. The best results are highlighted in bold. Dataset subclass names are replaced with initial capital letters.

Methods	BO	CA	CAP	CAR	GR	HA	LE	MN	PI	SC	TI	TO	TR	WO	ZI
CFA	0.068	0.056	0.050	0.271	0.004	0.341	<b>0.393</b>	0.255	0.080	0.015	0.155	0.053	0.056	0.281	0.573
DRAEM	0.117	0.019	0.044	0.018	0.005	0.036	0.013	0.142	0.104	0.002	0.130	0.039	0.040	0.033	<b>0.734</b>
PatchCore	0.048	0.029	0.035	0.552	0.003	0.338	0.279	0.248	0.051	0.008	0.249	0.034	0.079	0.304	0.595
RD4AD	0.055	0.040	0.064	0.212	0.005	0.384	0.116	0.247	0.061	0.015	0.193	0.034	0.059	0.097	0.562
SimpleNet	0.108	0.045	0.029	0.018	0.004	0.029	0.006	0.227	0.077	0.004	0.082	0.046	0.049	0.037	0.139
UniAD	0.054	0.031	0.022	0.047	0.007	0.189	0.053	0.110	0.034	0.008	0.107	0.040	0.045	0.103	0.444
PatchCore*	0.087	0.043	0.042	0.407	0.003	0.443	0.352	0.189	0.058	0.017	0.124	0.028	0.053	0.270	0.604
UniAD*	0.734	0.232	0.313	0.517	<b>0.204</b>	0.378	0.360	0.587	0.346	0.035	0.428	<b>0.398</b>	<b>0.542</b>	0.378	0.443
Ours	<b>0.752</b>	<b>0.290</b>	<b>0.349</b>	<b>0.622</b>	0.187	<b>0.506</b>	0.333	<b>0.775</b>	<b>0.634</b>	<b>0.214</b>	<b>0.549</b>	0.298	0.398	<b>0.535</b>	0.398

Table 2: Pixel-level AUPR $\uparrow$  on MVTec AD dataset (Bergmann et al. 2019) after training.

## Experiments and Discussion

### Experiments Setup

**Datasets** MVTec AD (Bergmann et al. 2019) is the most widely used dataset for industrial image anomaly detection. VisA (Zou et al. 2022) is now the largest dataset for real-world industrial anomaly detection with pixel-level annotations. We conduct experiments on these two datasets.

**Methods** We selected the most representative methods to establish the benchmark. These methods include CFA (Lee, Lee, and Song 2022), CSFlow (Rudolph et al. 2022), CutPaste (Li et al. 2021), DNE (Li et al. 2022), DRAEM (Zavrtanik, Kristan, and Skočaj 2021), FastFlow (Yu et al. 2021), FAVAE (Dehaene and Eline 2020), PaDiM (Defard et al. 2021), PatchCore (Roth et al. 2022), RD4AD (Deng and Li 2022), SPADE (Cohen and Hoshen 2020), STPM (Wang et al. 2021), SimpleNet (Liu et al. 2023c), and UniAD (You et al. 2022).

**Metrics** Following the common practice, we utilize Area Under the Receiver Operating Characteristics (AUROC/AUC) and Area Under Precision-Recall (AUPR/AP) for model evaluation. In addition, we use Forgetting Measure (FM) (Chaudhry et al. 2018) to evaluate models' ability to prevent catastrophic forgetting.

$$avg FM = \frac{1}{k-1} \sum_{j=1}^{k-1} \max_{l \in \{1, \dots, k-1\}} \mathbf{T}_{l,j} - \mathbf{T}_{k,j}, \quad (7)$$

where  $\mathbf{T}$  represents tasks,  $k$  stands for the current training task ID, and  $j$  refers to the task ID being evaluated. And  $avg FM$  represents the average forgetting measure of the model after completing  $k$  tasks. During the inference, we evaluate the model after training on all tasks.

**Training Details and Module Parameter Settings** We utilized the *vit-base-patch16-224* backbone pretrained on ImageNet 21K (Deng et al. 2009) for our method. During prompt training, we employed a batch size of 8 and adapt Adam optimizer (Kingma and Ba 2014) with a learning rate of 0.0005 and momentum of 0.9. The training process spanned 25 epochs. Our key-prompt-knowledge structure comprised a key of size (15, 196, 1024) float array, a prompt of size (15, 7, 768) float array, and knowledge of size (15, 196, 1024) float array, with an overall size of approximately 23.28MB.

### Continual Anomaly Detection Benchmark

We conducted comprehensive evaluations of the aforementioned 14 methods on the MVTec AD and VisA datasets.

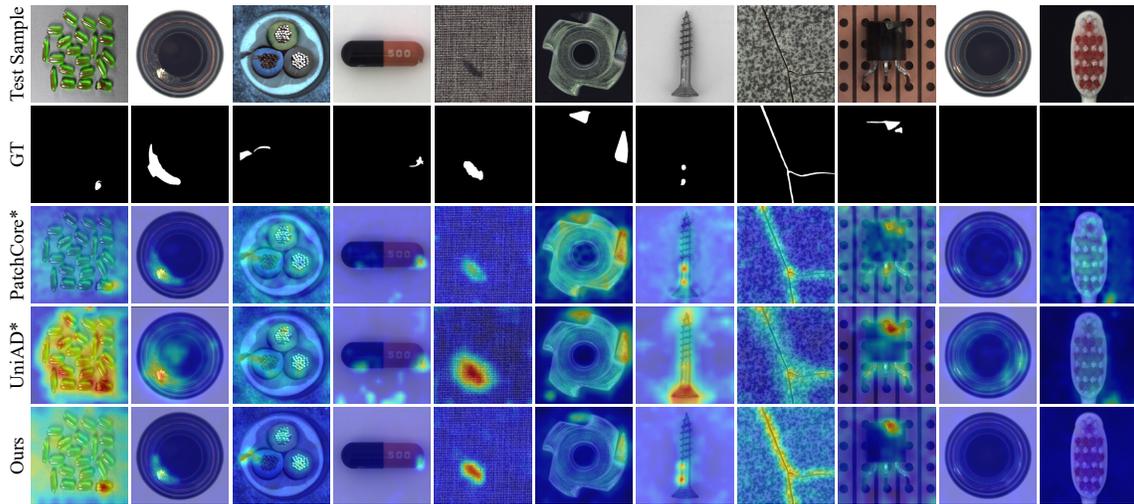


Figure 3: Visualization examples of continual anomaly detection. The first row displays the original anomaly images, the second row shows the ground truth annotations, and the third to fifth rows depict the heatmaps of our method and other methods.

Methods	Image AUROC		Pixel AUPR	
	Average	FM	Average	FM
CFA	0.623	0.361	0.177	0.083
CSFlow	0.539	0.426	-	-
CutPaste	0.312	0.510	-	-
DRAEM	0.595	0.371	0.098	0.116
FastFlow	0.512	0.279	0.044	0.214
FAVAE	0.547	0.102	0.083	0.083
PaDiM	0.545	0.368	0.086	0.366
PatchCore	0.602	0.383	0.190	0.371
RD4AD	0.596	0.393	0.143	0.425
SPADE	0.571	0.285	0.151	0.319
STPM	0.576	0.325	0.110	0.352
SimpleNet	0.708	0.211	0.060	0.069
UniAD	0.774	0.229	0.086	0.419
DNE	0.870	0.116	-	-
PatchCore*	0.669	0.318	0.181	0.343
UniAD*	0.904	0.076	0.393	0.086
Ours	0.930	0.010	0.456	0.013

Table 3: Image AUROC $\uparrow$ , Pixel-level AUPR $\uparrow$  and corresponding FM $\downarrow$  on MVTec AD dataset.

Among them, DNE stands as the SOTA method in unsupervised continual AD. Meanwhile, PatchCore and UniAD are two representative AD methods for memory-based and unified methods, respectively. Intuitively, these two methods appear to be better suited for the continual learning scenario. Due to the famous replay in continual learning methods, we also conducted replay-based experiments on PatchCore and UniAD. In these experiments, we provided them with a buffer capable of storing 100 training samples.

**Quantitative Analysis** As shown in Tables 1 - 5, most of the anomaly detection methods experienced significant performance degradation in the context of continual learning scenarios. Surprisingly, with the use of replay, UniAD managed to surpass DNE on the MVTec AD dataset. Moreover, on the VisA dataset, even without replay, UniAD out-

performed DNE. On the other hand, our method achieved a substantial lead over the second-best approach without the use of replay. Specifically, on the MVTec AD dataset, our method shows a 2.6 point lead in Image AUROC and a 6.3 point lead in Pixel AUPR over the second-ranked method, while on the VisA dataset, we achieve a 4.9 point lead in Image AUROC and a 1.7 point lead in Pixel AUPR. It can be observed that on the more complex structural VisA dataset, the detection capability of DNE, which solely relies on class tokens for anomaly discrimination, is significantly reduced. In contrast, our method remains unaffected.

Based on the comprehensive experimental results, our approach shows significant improvement over other methods in detecting anomalies under a continual setting. The experiments also demonstrate the potential of reconstruction-based methods, such as UniAD, in the field of continual UAD. In future works, combining our suggested CPM with the reconstruction-based UAD approach could be beneficial.

**Qualitative Analysis** As illustrated in Figure 3, our method demonstrates the ability to predict anomalies. This progress stands as a significant improvement compared to DNE. Compared to PatchCore\* and UniAD\*, our method exhibits two distinct advantages. Firstly, it demonstrates a more precise localization of anomalies. Secondly, it minimizes false positives in normal images.

### Ablation Study

**Module Effectivity** As shown in Table 6, We analyze the impact of two modules - Continual Prompting Module (CPM) and Structure-based Contrastive Learning (SCL). We observed significant improvements in the model’s performance with the implementation of these modules. In the absence of CPM’s key-prompt-knowledge architecture, our model used a single Knowledge base and reset it every time a new task was introduced. This approach restricted the model’s ability to adapt to continual learning without supervision. However, with the inclusion of CPM, the model’s

Methods	CA	CAP	CAS	CHW	FR	MA1	MA2	PCB1	PCB2	PCB3	PCB4	PF	AVG	FM
RD4AD	0.380	0.385	0.737	0.539	0.533	0.607	0.487	0.437	0.672	0.343	0.187	<b>0.999</b>	0.525	0.423
PatchCore	0.401	0.605	0.624	0.907	0.334	0.538	0.437	0.527	0.597	0.507	0.588	0.998	0.589	0.361
UniAD	0.573	0.599	0.661	0.758	0.504	0.559	0.644	0.749	0.523	0.547	0.562	0.989	0.639	0.297
DNE	0.486	0.413	0.735	0.585	0.691	0.584	0.546	0.633	0.693	0.642	0.562	0.747	0.610	0.179
PatchCore*	0.647	0.579	0.669	0.735	0.431	0.631	0.624	0.617	0.534	0.479	0.645	<b>0.999</b>	0.633	0.349
UniAD*	<b>0.884</b>	0.669	0.938	<b>0.970</b>	0.812	0.753	0.570	0.872	0.766	0.708	<b>0.967</b>	0.990	0.825	0.125
Ours	0.778	<b>0.877</b>	<b>0.960</b>	0.958	<b>0.945</b>	<b>0.823</b>	<b>0.667</b>	<b>0.905</b>	<b>0.871</b>	<b>0.813</b>	0.901	0.988	<b>0.874</b>	<b>0.039</b>

Table 4: Image-level AUROC $\uparrow$  and corresponding FM $\downarrow$  on VisA dataset (Zou et al. 2022) after training on the last subdataset.

Methods	CA	CAP	CAS	CHW	FR	MA1	MA2	PCB1	PCB2	PCB3	PCB4	PF	AVG	FM
RD4AD	0.002	0.005	0.061	0.045	0.098	0.001	0.001	0.013	0.008	0.008	0.013	0.576	0.069	0.201
PatchCore	0.012	0.007	0.055	0.315	0.082	0.000	0.000	0.008	0.004	0.007	0.010	0.585	0.090	0.311
UniAD	0.006	0.013	0.040	0.185	0.087	0.002	0.002	0.015	0.005	0.015	0.013	0.576	0.080	0.218
PatchCore*	0.018	0.010	0.047	0.202	0.081	0.003	0.001	0.008	0.004	0.008	0.010	0.443	0.070	0.327
UniAD*	<b>0.132</b>	0.123	0.378	<b>0.574</b>	<b>0.404</b>	<b>0.041</b>	<b>0.010</b>	0.612	0.083	<b>0.266</b>	<b>0.232</b>	0.549	0.283	0.062
Ours	0.067	<b>0.437</b>	<b>0.580</b>	0.503	0.334	0.013	0.003	<b>0.702</b>	<b>0.136</b>	<b>0.266</b>	0.106	0.457	<b>0.300</b>	<b>0.015</b>

Table 5: Pixel-level AUPR $\uparrow$  and corresponding FM $\downarrow$  on VisA dataset (Zou et al. 2022) after training on the last subdataset.

CPM	SCL	MVTec AD	VisA
$\times$	$\times$	0.693/0.183	0.584/0.050
$\checkmark$	$\times$	0.894/0.426	0.786/0.251
$\checkmark$	$\checkmark$	<b>0.930/0.456</b>	<b>0.874/0.300</b>

Table 6: Ablation study for CPM and SCL.

Encoder Layer	MVTec AD Metric	VisA Metric
1	0.840/0.399	0.806/0.143
3	0.934/0.451	<b>0.876/0.283</b>
5	0.930/ <b>0.456</b>	0.874/ <b>0.300</b>
7	<b>0.936/0.444</b>	0.872/0.267
9	0.906/0.420	0.853/0.248

Table 8: Ablation study for ViT encoder layer.

CPM	SCL	Knowledge Size	MVTec AD Metric	VisA Metric
$\checkmark$	$\times$	1x	0.894/0.426	0.786/0.251
$\checkmark$	$\times$	2x	0.921/0.452	0.818/0.255
$\checkmark$	$\times$	4x	0.929/0.453	0.860/0.294
$\checkmark$	$\checkmark$	1x	0.930/0.456	0.874/0.300
$\checkmark$	$\checkmark$	2x	0.936/0.461	0.893/0.307
$\checkmark$	$\checkmark$	4x	<b>0.938/0.466</b>	<b>0.909/0.310</b>

Table 7: Ablation study for *Knowledge* size and SCL.

Image AUROC score showed a significant improvement of 20 points. Regarding SCL, we found that without learning prompts contrastively, the model relied solely on the frozen ViT for feature extraction. This approach leads to a drop of around 4 points in the final performance, indicating the importance of SCL’s feature generalizability improvement.

**Size of Knowledge Base in CPM** To further investigate the role of SCL, we designed ablation experiments as illustrated in Table 7, by altering the size of Knowledge within the CPM module. The basic *Knowledge* size is 196, corresponding to the number of patches in a single image. Our method enables the representation of all images’ patches in a task with a single image feature space. Intriguingly, in the presence of SCL, even when the *Knowledge* size is 4 times larger, the performance enhancement remains marginal. However, without SCL, as the *Knowledge* size increases, the model exhibits a noticeable performance gain. This phenomenon can be attributed to SCL’s capacity to render feature distributions more compact, allowing a feature of the same size to encapsulate additional information.

**ViT Feature Layers** Furthermore, we explore the number of layers to use from ViT encoder in our method. The results in Table 8 indicate that neither shallow nor deep layers are effective for unsupervised anomaly detection. Intermediate layers, on the other hand, perform better as they are capable of representing both contextual and semantic information. We found that various datasets possess nuances in their definitions of anomalies, resulting in varying levels of granularity. While the degree of contextual knowledge required may vary across different datasets, we decided to stick with the fifth layer for simplicity.

## Conclusion

In this paper, we investigate the problem of applying continual learning on unsupervised anomaly detection to address real-world applications in industrial manufacturing. To facilitate this research, we build a comprehensive benchmark for unsupervised continual anomaly detection and segmentation. Furthermore, our proposed UCAD for task-agnostic CL on UAD is the first study to design a pixel-level continual anomaly segmentation method. UCAD novelty relies on a continual prompting module and structured-based contrastive learning, significantly improving continual anomaly detection performance. Comprehensive experiments have underscored our framework’s efficacy and robustness with varying hyperparameters. We also find that amalgamating and prompting ViT features from various layers might further enhance results, which we leave for future endeavors.

## Acknowledgments

This work is supported by the National Key R&D Program of China (Grant NO. 2022YFF1202903) and the National Natural Science Foundation of China (Grant NO. 62122035).

## References

- Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision*, 139–154.
- Aljundi, R.; Kelchtermans, K.; and Tuytelaars, T. 2019. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11254–11263.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD—A comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9592–9600.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2020. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4183–4192.
- Chandola, V.; Banerjee, A.; and Kumar, V. 2009. Anomaly Detection: A Survey. *ACM Computing Surveys*, vol. 41: 15.
- Chaudhry, A.; Dokania, P. K.; Ajanthan, T.; and Torr, P. H. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision*, 532–547.
- Cohen, N.; and Hoshen, Y. 2020. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*.
- Defard, T.; Setkov, A.; Loesch, A.; and Audigier, R. 2021. Padim: a patch distribution modeling framework for anomaly detection and localization. In *Proceedings of International Conference on Pattern Recognition*, 475–489. Springer.
- Dehaene, D.; and Eline, P. 2020. Anomaly localization by modeling perceptual features. *arXiv preprint arXiv:2008.05369*.
- Deng, H.; and Li, X. 2022. Anomaly Detection via Reverse Distillation from One-Class Embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9737–9746.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of 2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Eldar, Y.; Lindenbaum, M.; Porat, M.; and Zeevi, Y. Y. 1997. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9): 1305–1315.
- Gao, H.; Luo, H.; Shen, F.; and Zhang, Z. 2023. Towards Total Online Unsupervised Anomaly Detection and Localization in Industrial Vision. *arXiv preprint arXiv:2305.15652*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gudovskiy, D.; Ishizaka, S.; and Kozuka, K. 2022. Cflowad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 98–107.
- Hyun, J.; Kim, S.; Jeon, G.; Kim, S. H.; Bae, K.; and Kang, B. J. 2023. ReConPatch: Contrastive Patch Representation Learning for Industrial Anomaly Detection. *arXiv preprint arXiv:2305.16713*.
- Jiang, J.; Zhu, J.; Bilal, M.; Cui, Y.; Kumar, N.; Dou, R.; Su, F.; and Xu, X. 2022a. Masked Swin Transformer Unet for Industrial Anomaly Detection. *IEEE Transactions on Industrial Informatics*.
- Jiang, X.; Liu, J.; Wang, J.; Nie, Q.; Wu, K.; Liu, Y.; Wang, C.; and Zheng, F. 2022b. Softpatch: Unsupervised anomaly detection with noisy data. *Advances in Neural Information Processing Systems*, 35: 15433–15445.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *CoRR*, abs/1412.6980.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollar, P.; and Girshick, R. 2023. Segment Anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4015–4026.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. In *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Lee, S.; Lee, S.; and Song, B. C. 2022. Cfa: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization. *IEEE Access*, 10: 78446–78454.
- Lei, J.; Hu, X.; Wang, Y.; and Liu, D. 2023. Pyramid-Flow: High-Resolution Defect Contrastive Localization using Pyramid Normalizing Flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14143–14152.
- Li, C.-L.; Sohn, K.; Yoon, J.; and Pfister, T. 2021. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9664–9674.
- Li, W.; Gao, B.-B.; Xia, B.; Wang, J.; Liu, J.; Liu, Y.; Wang, C.; and Zheng, F. 2023. Cross-Modal Alternating Learning with Task-Aware Representations for Continual Learning. *IEEE Transactions on Multimedia*.
- Li, W.; Zhan, J.; Wang, J.; Xia, B.; Gao, B.-B.; Liu, J.; Wang, C.; and Zheng, F. 2022. Towards continual adaptation in industrial anomaly detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2871–2880.
- Liang, Y.; Zhang, J.; Zhao, S.; Wu, R.; Liu, Y.; and Pan, S. 2022. Omni-frequency channel-selection representations for unsupervised anomaly detection. *arXiv preprint arXiv:2203.00259*.

- Liu, J.; Xie, G.; ruitao chen; Li, X.; Wang, J.; Liu, Y.; Wang, C.; and Zheng, F. 2023a. Real3D-AD: A Dataset of Point Cloud Anomaly Detection. In *Proceedings of Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Liu, J.; Xie, G.; Wang, J.; Li, S.; Wang, C.; Zheng, F.; and Jin, Y. 2023b. Deep Industrial Image Anomaly Detection: A Survey. *arXiv preprint arXiv:2301.11514*, 2.
- Liu, X.; Ji, K.; Fu, Y.; Tam, W. L.; Du, Z.; Yang, Z.; and Tang, J. 2021. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*.
- Liu, Z.; Zhou, Y.; Xu, Y.; and Wang, Z. 2023c. Simplenet: A simple network for image anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20402–20411.
- Mishra, P.; Verk, R.; Fornasier, D.; Piciarelli, C.; and Foresti, G. L. 2021. VT-ADL: A vision transformer network for image anomaly detection and localization. *2021 IEEE 30th International Symposium on Industrial Electronics (ISIE)*, 01–06.
- Mousakhan, A.; Brox, T.; and Tayyub, J. 2023. Anomaly Detection with Conditioned Denoising Diffusion Models. *arXiv preprint arXiv:2305.15956*.
- Pirnay, J.; and Chai, K. 2022. Inpainting transformer for anomaly detection. In *Proceedings of International Conference on Image Analysis and Processing*, 394–406.
- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14318–14328.
- Rudolph, M.; Wandt, B.; and Rosenhahn, B. 2021. Same same but differnet: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1907–1916.
- Rudolph, M.; Wehrbein, T.; Rosenhahn, B.; and Wandt, B. 2022. Fully convolutional cross-scale-flows for image-based defect detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1088–1097.
- Salehi, M.; Sadjadi, N.; Baselizadeh, S.; Rohban, M. H.; and Rabiee, H. R. 2021. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14902–14912.
- Schlüter, H. M.; Tan, J.; Hou, B.; and Kainz, B. 2022. Natural synthetic anomalies for self-supervised anomaly detection and localization. In *Proceedings of European Conference on Computer Vision*, 474–489. Springer.
- Tien, T. D.; Nguyen, A. T.; Tran, N. H.; Huy, T. D.; Duong, S.; Nguyen, C. D. T.; and Truong, S. Q. 2023. Revisiting reverse distillation for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24511–24520.
- Wang, G.; Han, S.; Ding, E.; and Huang, D. 2021. Student-Teacher Feature Pyramid Matching for Anomaly Detection. *BMVC*.
- Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 139–149.
- Xie, G.; Wang, J.; Liu, J.; Jin, Y.; and Zheng, F. 2022. Pushing the Limits of Fewshot Anomaly Detection in Industry Vision: Graphcore. In *Proceedings of The Eleventh International Conference on Learning Representations*.
- Yan, X.; Zhang, H.; Xu, X.; Hu, X.; and Heng, P.-A. 2021. Learning semantic context from normal samples for unsupervised anomaly detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 3110–3118.
- You, Z.; Cui, L.; Shen, Y.; Yang, K.; Lu, X.; Zheng, Y.; and Le, X. 2022. A unified model for multi-class anomaly detection. *Advances in Neural Information Processing Systems*, 35: 4571–4584.
- Yu, J.; Zheng, Y.; Wang, X.; Li, W.; Wu, Y.; Zhao, R.; and Wu, L. 2021. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint arXiv:2111.07677*.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2021. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8330–8339.
- Zavrtanik, V.; Kristan, M.; and Skočaj, D. 2022. DSR—A dual subspace re-projection network for surface anomaly detection. *arXiv preprint arXiv:2208.01521*.
- Zhang, F.; and Chen, Z. 2023. IDDM: An incremental dual-network detection model for in-situ inspection of large-scale complex product. *Journal of Industrial Information Integration*, 33: 100463.
- Zhang, H.; Wang, Z.; Wu, Z.; and Jiang, Y.-G. 2023. DiffusionAD: Denoising Diffusion for Anomaly Detection. *arXiv preprint arXiv:2303.08730*.
- Zou, Y.; Jeong, J.; Pemula, L.; Zhang, D.; and Dabeer, O. 2022. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *Proceedings of European Conference on Computer Vision*, 392–408. Springer.