

Unveiling the Anomalies in an Ever-Changing World: A Benchmark for Pixel-Level Anomaly Detection in Continual Learning

Nikola Bugarin
University of Padova

nikola.bugarin@studenti.unipd.it

Jovana Bugarcic
University of Padova

jovana.bugaric@studenti.unipd.it

Manuel Barusco
University of Padova

manuel.barusco@phd.unipd.it

Davide Dalle Pezze
University of Padova

davide.dallepezze@unipd.it

Gian Antonio Susto
University of Padova

gianantonio.susto@unipd.it

Abstract

Anomaly Detection is a relevant problem in numerous real-world applications, especially when dealing with images. However, little attention has been paid to the issue of changes over time in the input data distribution, which may cause a significant decrease in performance. In this study, we investigate the problem of Pixel-Level Anomaly Detection in the Continual Learning setting, where new data arrives over time and the goal is to perform well on new and old data. We implement several state-of-the-art techniques to solve the Anomaly Detection problem in the classic setting and adapt them to work in the Continual Learning setting. To validate the approaches, we use a real-world dataset of images with pixel-based anomalies to provide a reliable benchmark and serve as a foundation for further advancements in the field. We provide a comprehensive analysis, discussing which Anomaly Detection methods and which families of approaches seem more suitable for the Continual Learning setting.

1. Introduction

Anomaly Detection (AD) is an important and challenging problem in the fields of Machine Learning and Computer Vision. Anomalies are patterns characterized by a noticeable deviation from the so-called normal data, where normal means compliant with some typical or expected features [44]. A significant advantage of the unsupervised techniques is that they do not require labeled data to learn from, making them easily deployable in many real-world applications. Indeed, in many applications, avoiding label collection is crucial, considering that this process is time-

expensive and resource-intensive, often requiring significant human effort and expertise. This becomes even more pronounced when our focus lies on details at the pixel-level, which can incur considerably higher costs than at the image-level [25]. Therefore, many approaches based on neural networks have been developed recently to address the problem of Unsupervised AD in the Computer Vision domain at the image-level and pixel-level [7, 49, 50]. Despite these advancements, significant challenges still prevent these approaches from being deployed. A significant issue is that shifts in the input data distribution can occur over time in real-world scenarios. For example, in the industrial sector, novel objects continually emerge, necessitating the identification of defects, while in the medical field, newly examined body structures require analysis to detect anomalies.

However, the neural networks are prone to the effect known as Catastrophic Forgetting: when learning a new task, artificial neural networks frequently forget the previous ones [42]. This aspect significantly hinders the adoption of these approaches in real-world scenarios. A new branch of Machine Learning has been introduced, known as Continual Learning (CL) [10] to address this issue. It focuses on learning from a stream of tasks to adapt to the new incoming tasks while remembering the previous ones.

Therefore, in this work, we study the use of Continual Learning strategies in the framework of Anomaly Detection on a data stream, focusing on Anomaly Localization, also known as Pixel-Level Anomaly Detection. We implement and test several well-known methods for AD that are considered state-of-the-art in the field. In particular, we adapt these AD techniques to work in the CL setting by employing the well-known CL strategy Replay [30, 36] when the methods allow it or perform ad hoc modifications of the original

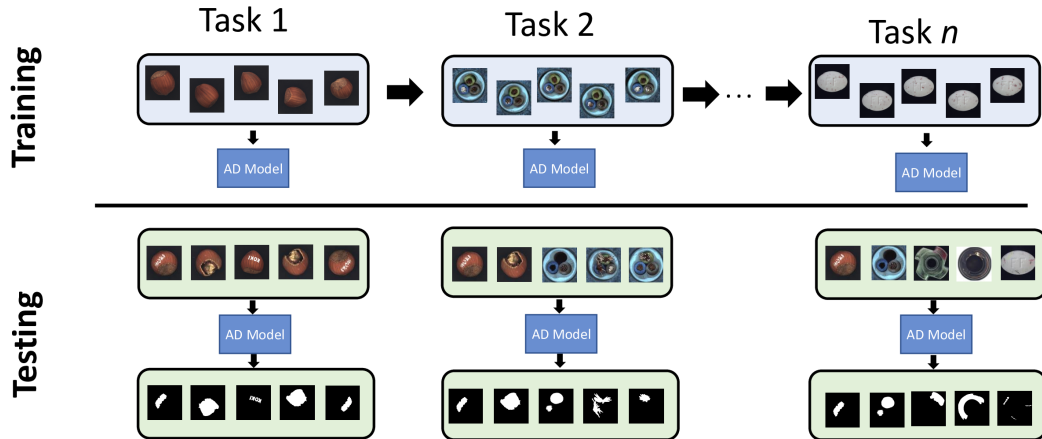


Figure 1. Considered CL setting for the AD problem. Each task corresponds to a new item. Our AD Model must be able to detect the anomalous products (image-level) and the defects inside the image (pixel-level) of a new item while remembering to perform well on previously seen items.

methods to let them work properly in the CL setting. Therefore, we perform a broad analysis by considering at least one approach for each category of the AD approaches in order to identify if some categories are more prone to forgetting than others. Then, we discuss which methods are more easily adaptable to the CL setting, which are more prone to forgetting, and which seem more fitting for the CL scenario. To summarize, this paper presents the following contributions:

- We explore the most well-known and state-of-the-art approaches for solving Pixel-Level Anomaly Detection within the Continual Learning setting.
- We adapt these AD techniques to work in the CL setting by employing the well-known CL strategy Replay when the methods allow it or performing ad hoc modifications of the original methods to enable them to work well.
- We provide a comprehensive analysis, discussing which AD methods and which families of approaches seem more suitable for the CL setting.

Moreover, to promote further research in the field and facilitate the rapid development of new methods along with comparisons with state-of-the-art approaches, we have made the code used available ¹. The paper is structured as follows: Sec. 2 provides an overview of the related work in the fields of CL and AD for Computer Vision. Sec. 4 details the AD methods implemented in our study and their adaptation for the CL setting. Sec. 5 presents the Experimental Setting, which includes information on the metrics used for evaluation and the benchmark dataset employed in our experiments. The Results of our experiments are presented in 6, where we analyze and discuss the performance of the adapted AD methods in the CL setting. Finally, in

7, we provide the Conclusion and discuss potential future research directions.

2. Related Work

We delineate the relevant literature as follows: Sec. 2.1 presents the CL framework. In Sec. 2.2, we outline the approaches for Unsupervised AD. Conversely, Sec. 2.3 delves into the intersection of CL and AD, highlighting previous research in this area.

2.1. Continual Learning

In the traditional Machine Learning setting, models are trained on fixed datasets. However, when considering real-world applications, it is easy to assume that the environments where the model is deployed could see new data with a different data distribution than the training distribution. CL addresses this by allowing DL models to update and expand their knowledge over time, with minimal computation and memory overhead [10]. Thus, an effective CL solution is expected to have low forgetting, require low memory consumption, and be computationally efficient.

The methods in the CL literature can be grouped into three big families of techniques known as rehearsal-based [3, 9, 26, 37, 40], regularization-based [18, 33, 35], and architecture-based approaches [28, 34, 41]. Rehearsal-based techniques assume storing and reusing past data samples during training. Within this category, numerous methodologies are available, yet one particularly distinguished approach is Experience Replay, also known as Replay [37]. Regularization-based approaches consider additional constraints or penalties during training to maintain the memory of old tasks [21, 24]. Architecture-based approaches, as the name suggests, focus on changing the orig-

¹<https://github.com/dallepezzo/clad>

inal model’s architecture to help maintain the old knowledge, and the methods to perform this differentiate significantly among them [12, 27, 39].

The related literature suggests that the Experience Replay approach appears to be the most effective and practical solution to reduce Catastrophic Forgetting [8, 20, 31, 47] when considering the image classification problem. However, when delving into CL strategies, it is essential to recognize that what works for one problem may not seamlessly translate to another. For instance, while rehearsal-based methods have shown remarkable efficacy in addressing issues like image classification, their applicability in domains such as Object Detection results could be limited, and distillation-based approaches like Lwf prove to be more effective [16].

Therefore, this highlights the importance of examining whether established methods, like Replay, are suitable for novel problems, such as Unsupervised AD. Moreover, the AD approaches vary greatly among them (they are split into several categories and sub-categories). This implies that a deep examination of how each AD technique performs based on a specific CL technique is necessary. Therefore, in this work, we provide a deep analysis of the use of the Replay approach in the problem of Pixel-Level Anomaly Detection, examining which methods are more suitable to work in a stream scenario. Moreover, as discussed below, while some AD techniques are well-suited to be used in the CL framework with the Replay approach, other methods are not adaptable to this technique. In these cases, we propose ad hoc modifications to let them operate in a continual scenario.

2.2. Anomaly Detection

Unsupervised AD approaches find many applications in Computer Vision (CV), encompassing manufacturing, the medical domain, autonomous vehicles, security systems, and more. In these contexts, detecting anomalies is crucial for safety, security, and efficiency. Many approaches proposed in the literature consider a model able to identify if a sample is anomalous. This information is then provided to users to help them in their decision-making process. However, they require more than just a binary outcome; they need an understanding of the reasons that brought the model to classify samples as anomalous. Without this understanding, decision-makers may struggle to formulate appropriate responses or interventions. Ensuring the interpretability of these systems can lead to safer and more efficient operations in various fields. This is achieved by the capability to extend the model’s predictions beyond image-level and into pixel-level granularity.

The advantage of providing an interpretable mechanism is not the only key advantage. Their unsupervised nature eliminates the need for a laborious label collection phase. This phase is typically time-consuming and resource-intensive,

demanding substantial human effort and expertise.

We can split most of the Pixel-Level AD approaches into two main families: reconstruction-based methods and Feature Embedding-based methods [4, 46].

Reconstruction-based methods learn to reconstruct normal images during training. The idea is to use generative models for data reconstruction, where a large error during reconstruction indicates the presence of an anomaly. This is the most historically significant research area, with approaches such as AutoEncoder and GANs [1, 6, 51, 52].

Feature Embedding-based methods consider data representations of images produced by a neural network, which is usually pre-trained. These approaches can be further categorized as Teacher-Student based, Normalizing Flow, and Memory Bank.

Teacher-Student approaches, as the name suggests, are based on two networks, a teacher and a student architecture. It exploits the knowledge distillation approach to transfer the learned knowledge, and when the features deviate, it is assumed the presence of an anomaly [5, 43].

Memory Bank approaches capture the features of normal images and store them in a feature memory bank [11, 22, 38]. Belonging to the category, three approaches are studied: Padim, PatchCore, and CFA (described with the rest of the methods in Sec. 4). All of these cannot be directly used for the Replay mechanism; therefore, ad hoc modifications are proposed to allow them to work effectively on a data stream. **Normalizing Flow approaches**, as implied by the name, are based on the normalizing flow models to transform the complex input data distribution into a normal distribution [17, 49]. Then, it becomes possible to use the probability of the data under the distribution as a measure of its normality.




Figure 2. Scheme of the classic Replay approach employed to solve the Anomaly Detection problem in the Continual Learning setting.

2.3. Continual Learning and Anomaly Detection

Even though AD is highly relevant in practice, only a few works in the literature address the problem of data distribution shifts for AD in the CL scenario. Most of the existing works consider tabular data or time series. For example,

[29] studied the industrial domain, while [19] considers the financial domain. In the domain of Network Intrusion Detection, a series of works have been proposed in the literature [44], [2], [15].

When considering the CLAD topic (Continual Learning for Anomaly Detection) in the field of Computer Vision, few works have been conducted. A study considering a sequence of tasks for AD was performed in [13] proposing a meta-learning approach, but they operate only at the image-level. Nevertheless, it is crucial to note that their methodology exclusively operates at the image-level and involves treating an entire class as anomalous, resulting in a scenario that diverges significantly from realistic situations.

A series of works are proposed that work at image-level. For example, [14] consider X-ray computed tomography images for a nuclear power plant. While they study and solve a realistic problem, their study exclusively examines image-level AD through supervised classification. On the contrary, [23] advocates for an unsupervised AD technique that retains samples from previous tasks. Despite moving from the supervised paradigm to the unsupervised one, this method exclusively operates at the image-level, diminishing its usefulness in real-world scenarios that demand interpretability. A comprehensive exploration of various AD techniques was performed in [45] by assessing multiple paradigms, including the CL setting. The study compared several AD methods, demonstrating how employing these methods using the Fine-Tuning strategy (without CL solutions) leads to forgetting. Notably, the study does not present any solutions for adapting these methods to function effectively within the CL setting. Finally, [32] studies a series of AD techniques in the CL setting, but most of the approaches are classic approaches and far from the state-of-the-art in the field.

Most of the work in the Computer Vision field focuses on predicting whether an image is normal or abnormal. However, Pixel-Level AD is frequently required in practice, given its several advantages, such as interpretability. Therefore, in our work, we implement several state-of-the-art AD techniques and make them work effectively in a data stream using the Replay approach when possible, or when the AD technique doesn't allow it, we propose some ad hoc modifications to make them work.

3. Continual Learning Approach

Among the CL strategies proposed in the literature, the Replay approach appears to be the most effective and practical solution to reduce Catastrophic Forgetting [8, 31, 47] when considering the image classification problem. Therefore, it is natural to consider it to be used to solve the AD problem in a data stream. The Replay strategy consists of storing some randomly selected samples from the previous tasks in raw format and using them to maintain knowledge about

previous tasks when training a new task. During training, each batch of data from the current task is combined with a batch of data from the memory of the previous tasks (sampled among them with the same probability). The entire process of the Replay approach is shown in Fig. 2.

It is crucial to note that this approach needs additional space for memorizing the subset of images from previous tasks. This implies that memory storage capacity is a major limitation; the model's performance decreases as memory capacity decreases. Replay serves as a highly versatile approach that, as we'll see, can be readily implemented across most Anomaly Detection (AD) techniques. Hence, unless specified otherwise, we employ the Replay approach for each AD technique, as depicted in Figure 2.

However, this approach is unsuitable to work (entirely or partially) when considering the Memory Bank approaches: Padim, PatchCore, and CFA. This is because these methods require a Memory Bank module that is not considered in the standard Replay approach. Therefore, for these approaches, specific ad hoc modifications are reported to let them work in a stream scenario. In other words, when discussing reconstruction-based approaches (DRAEM), Student-Teacher approaches like STFPM and EfficientAD, and distribution map-based approaches like FastFlow, Replay is implemented without particular challenges. Instead, when considering the Memory Bank approaches, ad hoc modifications are required as described below.

4. Anomaly Detection Methods

We tested all the following AD methods, briefly describing their characteristics and how they were adapted to work in the CL setting.

4.1. DRAEM

The Discriminatively Trained Reconstruction Embedding for Surface Anomaly Detection (DRAEM) method [52] is a reconstruction-based method. It is introduced to address a common drawback of generative AD methods. These methods only learn from anomaly-free data and lack optimization for discriminative AD since positive examples (anomalies) are unavailable during training. Training with synthetic anomalies leads to overfitting to synthetic appearances, resulting in poor generalization to real anomalies. To reduce overfitting, DRAEM proposes to train a discriminative model that considers the joint appearance of both reconstructed and original data, including the reconstruction subspace. Since it is purely based on learnable weights from neural network architectures, the classic Replay approach depicted in Fig. 2 is employed here.

4.2. STFPM

As the name suggests, Student-Teacher Feature Pyramid Matching (STFPM) [43] utilizes an architecture comprising two separate models for student and teacher, respectively. In contrast to typical Knowledge Distillation methods, here the teacher and student networks in STFPM share the same architecture. The distinction is that the teacher network is pre-trained on Imagenet dataset, while the student network begins with random initialization. During the training phase, the objective is to enable the student network to reproduce the feature maps generated by the teacher. During the test phase, differences across different network levels are computed to produce the anomaly map of an image.

4.3. EfficientAD

EfficientAD [5] employs a Student-Teacher methodology similar to STFPM. However, it introduces a novel approach for feature extraction from a pre-trained neural network using a significantly reduced-depth model called Patch Description Network (PDN). Notably, Efficient AD utilizes a single teacher and a single student network based on the PDN architecture. Furthermore, to address logical anomalies such as misplaced objects, an autoencoder is integrated into the EfficientAD architecture. As for STPFM, the classic Replay approach is employed.

4.4. Padim

Patch Distribution Modeling (PaDiM) is introduced by [11]. Similar to other feature embedding-based techniques, it leverages a pre-trained CNN to extract patch embeddings. They propose that each position of a patch in the image can be characterized by a multivariate Gaussian distribution. During the inference process for a given test image, the Mahalanobis distance is computed for each patch, providing the anomaly score. In this approach, there are no weights from a neural network to be updated, only a set of multivariate Gaussian distributions, one for each patch. Typically, this implies considering a different set of Gaussian distributions for each task, but this violates constraints on the memory size that should not increase linearly with the number of tasks.

Therefore, to allow the Padim approach to be updated with constant memory size, we propose an incremental averaging of Gaussian parameters when introducing new tasks (see the Supplementary Material ² for more details). The sum of old and new Gaussian distributions allows the memory to remain constant over time while allowing the model to have a representation of all tasks seen so far using a single Gaussian distribution per patch.

²Supplementary Material can be found at <https://github.com/dallepezz/clad>

4.5. PatchCore

PatchCore [38] is a Memory Bank-based approach where patch features are stored within a memory bank. However, to enhance efficiency, a coreset-reduction approach is employed to reduce the number of patches to be stored. In this approach, there are no weights from a neural network to be updated, only the memory bank that needs to be addressed.

Upon the encounter of a new task, the coreset-reduction algorithm is utilized to identify m patches to be retained, where $m = \frac{1}{N} \text{memory size}$, with N denoting the number of tasks seen so far. The newly identified m patches are incorporated into the memory bank. At the same time, the number of patches associated with the previous tasks is reduced to maintain only m patches for each old task. To achieve this, the patches associated with each task undergo coreset-reduction until m patches are attained. (More details in the Supplementary Material ²)

4.6. CFA

In [22], a novel Memory Bank approach called Coupled-Hypersphere-Based Feature Adaptation (CFA) is proposed. Initially, a memory bank is created based on the features extracted from the training set. Then, a trainable Patch Descriptor Network (PDN) is proposed with the aim of increasing the density of normal features (in alignment with the content stored in the memory bank). The intuition is that this will allow us to distinguish the abnormal features more clearly.

In this method, two distinct sources are identified that may contribute to the phenomenon of forgetting, the memory bank and the PDN network. To address the memory bank problem, we propose to update the memorized patch features using an incremental average across batches of the current task training dataset (further elaboration provided in the Supplementary Material). After completing this initial procedure upon receiving the new task, we proceed to update the network parameters of the Patch Descriptor. To perform the update of this part, we consider the Replay approach, consistent with its application in previous methods like DRAEM, STFPM, and EfficientAD. The combination of these two mechanisms enables the model to learn from a stream of tasks while maintaining a constant memory size over time.

4.7. FastFlow

FastFlow [49] utilizes 2D normalizing flows as a probability distribution estimator. It takes as input the features obtained from a deep feature extractor. During training, FastFlow learns to transform the input visual features into a tractable distribution, which is used to recognize anomalies in the inference phase. This model employs the Replay approach as described in Sec. 3 and as represented in Fig. 2.

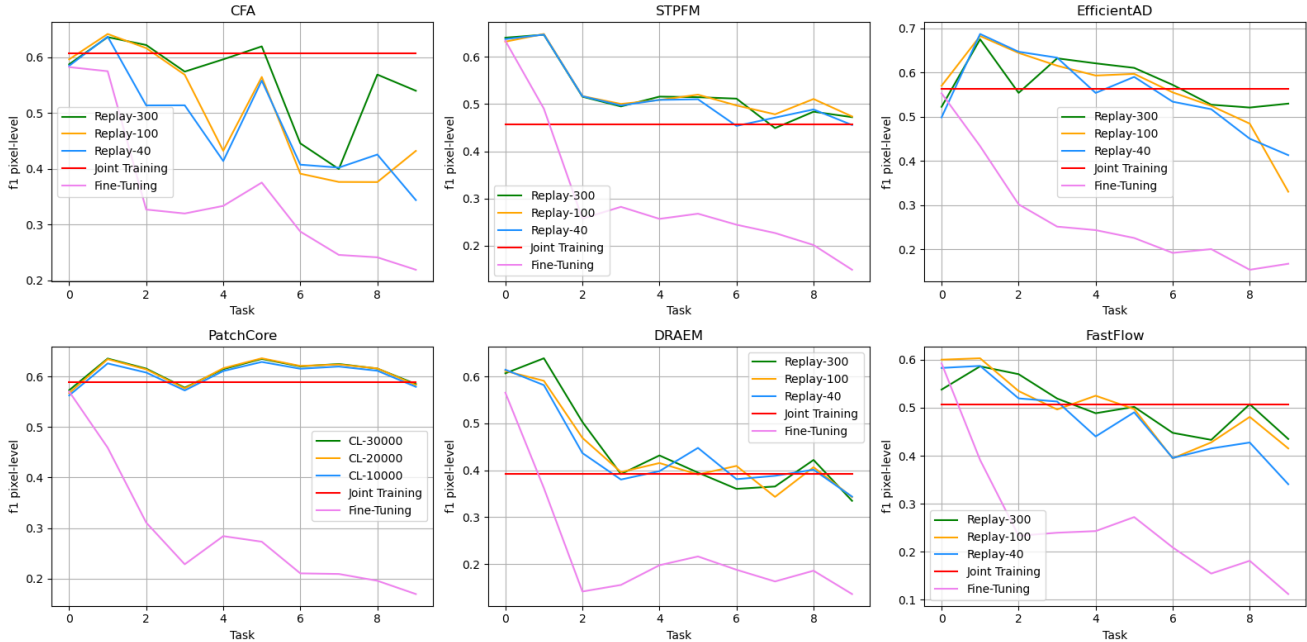


Figure 3. Each plot shows an AD technique tested in the CL setting. In detail, each plot shows the behavior of the f1 pixel-level metric over time, where each point represents the average performance on all the tasks seen so far. The considered strategies are the same as defined in Sec. 5.

5. Experimental Setting

In the following part, we delineate the experimental setup employed in our conducted studies. Specifically, we outline the CL scenario in Sec. 5.1. Following this, in Sec. 5.2, we explain the metrics under consideration, which are shown in Tab. 1. Lastly, we describe the methods present in Fig. 3 used as comparison with the CL approaches.

5.1. Benchmark Dataset

The MVTEC Dataset [6] is regarded as an extensive repository of images designed specifically for evaluating AD algorithms. This dataset encompasses ten objects and five textures, making it suitable for assessing the robustness and generalization capabilities of various AD techniques. In the given CL scenario, we consider a sequence of ten tasks, where each task corresponds to a different object (similarly as depicted in Fig. 1). Specifically, the evaluation includes the following objects: Bottle, Cable, Capsule, Hazelnut, Transistor, Metal Nut, Pill, Screw, Zipper, and Toothbrush.

5.2. Metrics

5.2.1 Anomaly Detection metrics

Various evaluation metrics are commonly employed to assess the performance of AD techniques using the MVTEC dataset. The first important categorization is based on how the evaluation is performed, whether on the image or pixel

level. For both of them, we consider the ROC and f1 metrics, as reported in Tab. 1. For the pixel-level, we also consider the Per-region-overlap (PRO) metric, which ensures equal weighting of ground-truth regions regardless of their sizes, as explained in [48]. This stands in stark contrast to simplistic per-pixel metrics, where a single large correctly segmented region can compensate for numerous inaccurately segmented smaller ones.

5.2.2 Continual Learning metrics

The metrics described earlier are proposed within the CL setting. For each metric listed in Tab. 1, the final value is averaged across the set of tasks, as it is standard practice in CL evaluations, such as accuracy metric. Furthermore, the percentage of forgetting is reported to quantify the extent of forgetting across tasks. Additionally, to provide a more comprehensive assessment, the relative gap between the CL approach and the Joint Training strategy is highlighted. This comparison is essential as a model may exhibit low forgetting but substantially lower performance compared to the Joint Training strategy, which is the upper bound for all the CL approaches.

As previously stated, in the CL framework we are interested in updating and expanding their knowledge over time, avoiding forgetting, while at the same time we want to obtain this goal with minimal computation and memory over-

Table 1. Performance comparison of Anomaly Detection strategies in the Continual Learning setting. Green reports the best value for each metric, while red shows the worst. All the metrics considered are described in Sec. 5, while each tested AD method is explained in 4.

Strategy		Reconstruction based	Memory Bank			Student Teacher		Normalizing Flow
		DRÆM	PatchCore	PaDiM	CFA	STPFM	EfficientAD	FastFlow
Performance		<i>Replay</i>	<i>CL</i>	<i>CL</i>	<i>Replay</i>	<i>Replay</i>	<i>Replay</i>	<i>Replay</i>
Image - level	<i>AUC ROC</i>	0.75	0.97	0.56	0.92	0.87	0.81	0.84
	<i>f1</i>	0.87	0.97	0.83	0.94	0.89	0.89	0.89
Pixel - level	<i>AUC ROC</i>	0.82	0.98	0.73	0.93	0.93	0.88	0.91
	<i>f1</i>	0.32	0.58	0.17	0.53	0.48	0.49	0.42
	<i>PR AUC</i>	0.11	0.55	0.29	0.49	0.43	0.43	0.35
	<i>AU PRO</i>	0.65	0.89	0.53	0.75	0.85	0.67	0.73
Time	<i>training</i>	11h 46min	8min	6min	2h 16min	27min	5h 26min	2h 19min
Architecture memory [MB]		389.6	275.6	275.6	301.2	93.6	82.8	440.80
Additional memory [MB]		59.0	184.3	1541.0	67.7	59.0	59.0	59.0
Relative gap (δ) [%]		7.50	0.01	35.00	7.33	2.00	7.00	9.00
Average forgetting [%]		27.08	0.72	19.12	-4.70	14.78	19.11	17.39

head [10]. Hence, the values for memory and computation are also provided in Tab. 1. For memory, we report the Architecture Memory and Additional Memory, indicating with Total Memory their combination. With Additional Memory we refer to the memory necessary to store images for the Replay approach of patches for the Memory Bank based approaches. Instead, Training Time is provided to assess computation.

5.3. Considered CL Strategies

During our work, we considered the comparison among the following CL approaches:

- **Joint Training:** The model is trained on all tasks simultaneously, which is usually considered as an upper baseline for the CL strategies since they are not affected by the phenomenon of Catastrophic Forgetting.
- **Fine-tuning:** As a lower bound, we consider the fine-tuning approach, in which a model is presented sequentially only with data from the current task.
- **Replay:** The approach as described in Sec. 3. Specifically, we consider 300, 100, and 40 images that can be stored in the Replay memory, representing respectively 10%, 5%, and 2% of the entire dataset.
- **CL:** We report this generic signature for all the memory-bank approaches like Padim and PatchCore that do not fit the Replay category and have required ad hoc modifications to perform well in a data stream. In particular, for PatchCore, we vary the memory bank size to 30000, 20000, 10000.

6. Results

This section analyzes the results outlined in Tab. 1 and Fig. 3. For each AD method, a different plot is shown in Fig.

3, and a comprehensive comparison for each method and metric is also provided in Tab. 1. We specifically examine the AD performance, memory consumption, and training time in Sec. 6.1, 6.2, and 6.3 respectively. Subsequently, in Sec. 6.4 a generic analysis is provided for each method and AD families of techniques when considering all three factors together.

6.1. Analysis on the performance metrics

A general observation, as notable from Fig. 3, is that for all AD methods (except Padim) the degree of forgetting is not particularly pronounced. The obtained performance resembles the ones of the upper bound Joint Training. Notably, the PatchCore method demonstrates no sign of forgetting, which is likely attributed to its reliance on a memory bank. Of particular interest is the surprising performance of the CFA method, which appears to have improved instead of experiencing a decline in performance. This phenomenon could be attributed to the methodology of the CFA approach, which facilitates the gradual convergence of features from old and new tasks while effectively isolating anomalies.

In contrast, all methods reliant on large networks with learnable parameters show some level of forgetting, typically falling within the range of 14.78% to 19.11%, except for DRAEM, which exhibits a higher forgetting rate of 27.08%. Notably, for most approaches, the gap concerning the upper bound is low and around 7%. Consequently, the Replay approach effectively addresses the AD problem within the CL framework. In greater detail, the PatchCore method emerges as the optimal choice, achieving a f1 pixel-level score of 0.58. Good performance is also shown in CFA, STPFM, and Efficient AD (0.54, 0.49, 0.48, respectively).

The lower performance is encountered with FastFlow (0.42) and DRAEM (0.32).

6.2. Analysis on the memory consumption

As discussed before, it is crucial to consider memory consumption when evaluating methods. On this front, we can observe that PaDiM is the least efficient technique for analyzing memory consumption, consuming a total memory of 1816MB. Following are the methods PatchCore, CFA, DRAEM, and FastFlow that have medium memory usage falling in the 369MB - 500MB range. Then, the clear winners are the student teacher-based methods, i.e., EfficientAD and STFPM, which require only 142MB and 153MB, respectively. These observations are made when considering the Total Memory, i.e., the combination of Architecture Memory and Additional Memory.

Focusing on the Additional Memory component, which preserves data representations in methods like Replay or Memory Bank, can yield some interesting insights. Despite the significant memory consumption, Padim continues to exhibit poor performance, utilizing 1541MB. In contrast, the method boasting the highest f1 pixel-level score, PatchCore, also demands the most memory, totaling 184.3MB. Notably, CFA follows with 67.7MB, followed by all Replay-based approaches (STFPM, EfficientAD, DRAEM, FastFlow), each requiring 59MB of Replay memory. Essentially, Memory Bank-based approaches tend to require more memory than student-teacher-based, normalizing-flow-based, and reconstruction-based methods.

6.3. Analysis on the Training Time

When considering only the Training Time, Padim shows the fastest training time but consistently shows the poorest performance in terms of AD and Total Memory. Then Patchcore shows a training time of just 8 minutes. Another noteworthy method, STFPM, distinguishes itself with its rapid training duration of only 27 minutes. This contrasts with EfficientAD, another student-teacher-based approach, which requires considerably longer training time, totaling 5 hours and 26 minutes. Additionally, FastFlow and CFA demonstrate relatively good training times, each taking approximately 2 hours. Conversely, DRAEM emerges as the slowest option, necessitating 11 hours and 46 minutes.

6.4. Comprehensive Evaluation

In this part, we provide a comparison, considering in conjunction all three critical factors: AD Performance, Memory, and Training Time. The optimal choice depends on the specific requirements defined by the user. For instance, we can assess PatchCore as the main winner when considering performance and training time. However, other approaches could be adopted when including additional considerations, such as Total Memory, like the STFPM technique, which

shows a good trade-off among all factors evaluated. When considering the families of approaches in general, the Memory Bank-based approaches exhibit excellent AD performance and training efficiency but slightly higher memory consumption. In contrast, student-teacher-based, normalizing flow-based, and reconstruction-based approaches seem to be good choices when the memory consumed is critical. Nevertheless, the reconstruction-based approaches do not offer a good choice regarding AD and training time.

7. Conclusions and Future Work

In this work, we studied the Continual Learning Setting for the Pixel-Level Anomaly Detection problem. This scenario is of paramount importance since traditional AD methods often operate under the assumption of stationary data distributions, which is not realistic in many practical scenarios. In our study, we effectively integrate several state-of-the-art AD techniques into the CL framework. To make these AD techniques operate effectively in the CL setting, we propose using the well-known Replay approach to make them operate in a data stream or perform ad hoc modifications when Replay is not possible. In particular, we provide insights into each method's strengths and weaknesses, highlighting its performance, memory consumption, training time, and resilience to forgetting.

An interesting future research direction is the evaluation of additional AD techniques not considered in this study, which could provide a different insight concerning the approaches considered in this work. Furthermore, it would be of extreme relevance an investigation of CL approaches that are not rehearsal-based. Indeed, while Replay proved a viable and effective solution to solve the AD problem in the CL setting, we do not know the behavior of other approaches like distillation-based approaches (e.g., LwF) or methods like EWC. Moreover, as pointed out in our study, future studies should focus on methods that try to optimize AD performance, memory, and training time simultaneously to guarantee deployment feasibility when considering resource-constrained environments.

Acknowledgments

This study was partially carried out within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from Next-GenerationEU (Italian PNRR – M4 C2, Invest 1.3 – D.D. 1551.11-10- 2022, PE00000004).

References

- [1] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia*,

- December 2–6, 2018, Revised Selected Papers, Part III 14, pages 622–637. Springer, 2019. 3
- [2] Suresh Kumar Amalapuram, Akash Tadwai, Reethu Vinta, Sumohana S Channappayya, and Bheemarjuna Reddy Tamma. Continual learning for anomaly based network intrusion detection. In *2022 14th International Conference on COMMunication Systems & NETWORKS (COMSNETS)*, pages 497–505. IEEE, 2022. 4
- [3] Craig Atkinson, Brendan McCane, Lech Szymanski, and Anthony Robins. Pseudo-rehearsal: Achieving deep reinforcement learning without catastrophic forgetting. *Neuro-computing*, 428:291–307, 2021. 2
- [4] Jinan Bao, Hanshi Sun, Hanqiu Deng, Yinsheng He, Zhaoxiang Zhang, and Xingyu Li. Bmad: Benchmarks for medical anomaly detection. *arXiv preprint arXiv:2306.11876*, 2023. 3
- [5] Kilian Batzner, Lars Heckler, and Rebecca König. Efficientad: Accurate visual anomaly detection at millisecond-level latencies. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 128–138, 2024. 3, 5
- [6] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9584–9592, 2019. 3, 6
- [7] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9584–9592, 2019. 1
- [8] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2180–2187. IEEE, 2021. 3, 4
- [9] Xi Chen, Jie Jiang, Zhiqiang Li, Honggang Qi, Qingli Li, Jiapeng Liu, Laiwen Zheng, Min Liu, and Yongqiang Deng. An online continual object detector on vhr remote sensing images with class imbalance. *Engineering Applications of Artificial Intelligence*, 117:105549, 2023. 2
- [10] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. 1, 2, 7
- [11] Thomas Defard, Aleksandr Setkov, Angélique Loesch, and Romaric Audigier. PaDiM: A patch distribution modeling framework for anomaly detection and localization. In *Pattern Recognition. ICPR International Workshops and Challenges*, pages 475–489. Springer International Publishing, 2021. 3, 5
- [12] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. 3
- [13] Ahmed Frikha, Denis Krompaß, and Volker Tresp. Arcade: A rapid continual anomaly detector. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10449–10456. IEEE, 2021. 4
- [14] Hossam A Gabbar, Oluwabukola Grace Adegboro, Abderazak Chahid, and Jing Ren. Incremental learning-based algorithm for anomaly detection using computed tomography data. *Computation*, 11(7):139, 2023. 4
- [15] Gastón García González, Pedro Casas, Alicia Fernández, and Gabriel Gómez. Steps towards continual learning in multivariate time-series anomaly detection using variational autoencoders. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 774–775, 2022. 4
- [16] Linting Guan, Yan Wu, Junqiao Zhao, and Chen Ye. Learn to detect objects incrementally. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 403–408. IEEE, 2018. 3
- [17] Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 98–107, 2022. 3
- [18] Ya-nan Han and Jian-wei Liu. Online continual learning via the meta-learning update with multi-scale knowledge distillation and data augmentation. *Engineering Applications of Artificial Intelligence*, 113:104966, 2022. 2
- [19] Hamed Hemati, Marco Schreyer, and Damian Borth. Continual learning for unsupervised anomaly detection in continuous auditing of financial accounting data. *arXiv preprint arXiv:2112.13215*, 2021. 4
- [20] Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. Imbalanced continual learning with partitioning reservoir sampling. In *European Conference on Computer Vision*, pages 411–428. Springer, 2020. 3
- [21] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 2
- [22] Sungwook Lee, Seunghyun Lee, and Byung Cheol Song. Cfa: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization. *IEEE Access*, 10:78446–78454, 2022. 3, 5
- [23] Wujin Li, Jiawei Zhan, Jinbao Wang, Bizhong Xia, Bin-Bin Gao, Jun Liu, Chengjie Wang, and Feng Zheng. Towards continual adaptation in industrial anomaly detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2871–2880, 2022. 4
- [24] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 2
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1

- [26] Yongle Luo, Yuxin Wang, Kun Dong, Qiang Zhang, Erkang Cheng, Zhiyong Sun, and Bo Song. Relay hindsight experience replay: Self-guided continual reinforcement learning for sequential object manipulation tasks with sparse rewards. *Neurocomputing*, 557:126620, 2023. 2
- [27] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 3
- [28] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018. 2
- [29] Benjamin Maschler, Thi Thu Huong Pham, and Michael Weyrich. Regularization-based continual learning for anomaly detection in discrete manufacturing. *Procedia CIRP*, 104:452–457, 2021. 4
- [30] Gabriele Merlin, Vincenzo Lomonaco, Andrea Cossu, Antonio Carta, and Davide Bacciu. Practical recommendations for replay-based continual learning methods. 1
- [31] Lorenzo Pellegrini, Gabriele Graffieti, Vincenzo Lomonaco, and Davide Maltoni. Latent replay for real-time continual learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10203–10209. IEEE, 2020. 3, 4
- [32] Davide Dalle Pezze, Eugenia Anello, Chiara Masiero, and Gian Antonio Susto. Continual learning approaches for anomaly detection. *arXiv preprint arXiv:2212.11192*, 2022. 4
- [33] Jary Pomponi, Simone Scardapane, Vincenzo Lomonaco, and Aurelio Uncini. Efficient continual learning in neural networks with embedding regularization. *Neurocomputing*, 397:139–148, 2020. 2
- [34] Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [35] Jason Ramapuram, Magda Gregorova, and Alexandros Kalousis. Lifelong generative modeling. *Neurocomputing*, 404:381–400, 2020. 2
- [36] ANTHONY ROBINS. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. 1
- [37] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [38] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. *arXiv:2106.08265*, 2022. 3, 5
- [39] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 3
- [40] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017. 2
- [41] Ghada Sokar, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Spacenet: Make free space for continual learning. *Neurocomputing*, 439:1–11, 2021. 2
- [42] Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. *CoRR*, abs/1904.07734, 2019. 1
- [43] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. *arXiv:2103.04257*, 2021. 3, 5
- [44] Felix Wiewel and Bin Yang. Continual learning for anomaly detection with variational autoencoder. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3837–3841, 2019. 1, 4
- [45] Guoyang Xie, Jinbao Wang, Jiaqi Liu, Jiayi Lyu, Yong Liu, Chengjie Wang, Feng Zheng, and Yaochu Jin. Im-iad: Industrial image anomaly detection benchmark in manufacturing. *arXiv preprint arXiv:2301.13359*, 2023. 4
- [46] Guoyang Xie, Jinbao Wang, Jiaqi Liu, Jiayi Lyu, Yong Liu, Chengjie Wang, Feng Zheng, and Yaochu Jin. Im-iad: Industrial image anomaly detection benchmark in manufacturing. *IEEE Transactions on Cybernetics*, 2024. 3
- [47] Qihan Yang, Fan Feng, and Rosa Chan. A benchmark and empirical analysis for replay strategies in continual learning, 2022. 3, 4
- [48] Yifei Yang, Shibing Xiang, and Ruixiang Zhang. Improving unsupervised anomaly localization by applying multi-scale memories to autoencoders. *arXiv preprint arXiv:2012.11113*, 2020. 6
- [49] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows, 2021. 1, 3, 5
- [50] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021. 1
- [51] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021. 3
- [52] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. DRAEM - A discriminatively trained reconstruction embedding for surface anomaly detection. *arXiv:2108.07610*, 2021. 3, 4