

Multi-Label Continual Learning for the Medical Domain: A Novel Benchmark

Marina Ceccon*, Davide Dalle Pezze*, Alessandro Fabris**, Gian Antonio Susto*

*University of Padova, Padova, Italy

**Max Planck Institute for Security and Privacy, Bochum, Germany

marina.ceccon@phd.unipd.it, davide.dallepezze@unipd.it,
 alessandro.fabris@mpi-sp.org, gianantonio.susto@unipd.it

Abstract

Despite the critical importance of the medical domain in Deep Learning, most of the research in this area solely focuses on training models in static environments. It is only in recent years that research has begun to address dynamic environments and tackle the Catastrophic Forgetting problem through Continual Learning (CL) techniques. Previous studies have primarily focused on scenarios such as Domain Incremental Learning and Class Incremental Learning, which do not fully capture the complexity of real-world applications. Therefore, in this work, we propose a novel benchmark combining the challenges of new class arrivals and domain shifts in a single framework, by considering the New Instances and New Classes (NIC) scenario. This benchmark aims to model a realistic CL setting for the multi-label classification problem in medical imaging. Additionally, it encompasses a greater number of tasks compared to previously tested scenarios. Specifically, our benchmark consists of two datasets (NIH and CXP), nineteen classes, and seven tasks. To solve common challenges (e.g., the task inference problem) found in the CIL and NIC scenarios, we propose a novel approach called Replay Consolidation with Label Propagation (RCLP). Our method surpasses existing approaches, exhibiting superior performance with minimal forgetting.

1. Introduction

In recent years, several studies have proven the efficacy of using Deep Learning models to detect diseases from chest X-ray images [26, 51]. While these models allowed for improved state-of-the-art in the medical field, several significant challenges must be addressed before considering their use to support decision-making in realistic scenarios. One such challenge is ensuring the model's adaptability in dynamic environments, where shifts in input data distribution may occur over time [25, 42].

For instance, each hospital may employ machines with

different image acquisition techniques; this can perturb the performance of a previously trained model, requiring ongoing recalibration and refinement [26]. Moreover, the need to detect diseases not initially incorporated into the trained model may arise. For instance, the labeling of additional classes may have been provided at a later point in time, reflecting evolving medical understanding and diagnostic criteria [40].

Failure to address these challenges can severely limit the model's utility and constrain its ability to effectively serve patient needs. A paradigm called Continual Learning (CL) has emerged in the literature to solve this problem. CL enables the model to adapt to new data while retaining knowledge from the old data. Previous studies in this domain focused on solving the problem in two different settings: Domain Incremental Learning (DIL) and Class Incremental Learning (CIL). However, these scenarios fail to accurately represent real-world challenges where both new classes and domain shifts can happen. Therefore, we propose a novel benchmark for the medical imaging field based on the New Instances & New Classes (NIC) scenario [30]. Specifically, we evaluate this scenario in the context of pathology classification of chest X-ray images, considering nineteen classes, seven tasks, and two domains. This setting combines the challenges of both new class arrivals and domain shifts within a single framework, mirroring the complexities often encountered in realistic applications like medical imaging [30].

On the proposed benchmark, the most well-known techniques in the CL field are tested. In particular, replay-based and distillation-based approaches are tested since they are often used to solve the CIL scenario in image classification [25, 40]. However, these approaches face limitations in multi-label settings such as the one considered. Replay-based approaches may suffer from task interference between the samples from the current task and the replayed samples [44]. Instead, distillation-based methods deteriorate when old classes do not reappear in future tasks [19]. Therefore, we introduce a novel approach called Re-

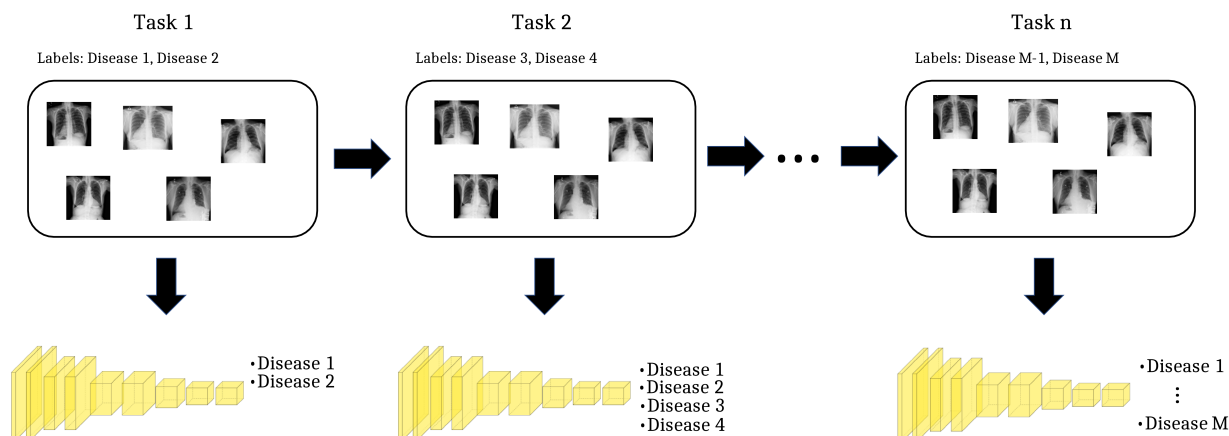


Figure 1. Scheme of the multi-label CL setting in the context of classification of chest X-rays. Diagnostic capabilities are expanded over time with new tasks.

play Consolidation with Label Propagation (RCLP) to solve these challenges.

Overall, we make the following contributions:

- We introduce a novel benchmark for CL in medical imaging, combining the challenges of new class arrivals and domain shifts in a single framework.
- We propose Replay Consolidation with Label Propagation, a novel method to address the multi-label image classification problem in the medical imaging setting.
- Experimental results demonstrate the effectiveness of RCLP, outperforming existing methods and achieving performance improvements.

In addition, to promote further research in the domain and facilitate the advancement of novel methodologies along with comparisons with state-of-the-art approaches, we make the code available.¹

Our work is structured as follows. Sec. 2 provides an overview of the CL literature with a focus on the medical domain. Sec. 3 presents the proposed benchmark for multi-label medical imaging, while Sec. 4 proposes the RCLP approach. In Sec. 5 describes the experimental setting, and in Sec. 6 reports the results of our experiments. Lastly, Sec. 7 concludes this work by discussing limitations and future research directions.

2. Related Work

In Sec. 2.1, we provide an overview of the most popular CL scenarios and methods. Following that, we focus on the

medical domain and chest X-ray image classification for the CL setting in Sec. 2.2.

2.1. Continual Learning

CL has emerged as a framework to adapt models to new data distributions without forgetting what they have already learned [27]. The main problem tackled by the CL methods is the so-called catastrophic forgetting. By fine-tuning the new distributions of data, the knowledge of old data gets completely corrupted [23, 35]. In the CL literature, most works refer to three CL scenarios: Domain-Incremental, Class-Incremental, and Task-Incremental Learning [45]. In the DIL scenario, each new task presents a shift in the input data distribution but the classes are always the same [27]. In the CIL scenario, each task contains new classes. The Task-Incremental scenario, contrary to the previous two, has the peculiarity of providing the task ID during the testing phase.

The methods in the CL literature can be grouped into three big families. *Rehearsal-based* techniques assume storing and reusing past data samples during training. Within this category, various approaches exist; a prominent method is Experience Replay [36], simply referred to as *Replay*. This method stores a portion of the old data in memory. While training on a new task, the new data is combined with the data from the memory buffer to maintain knowledge of old tasks.

Regularization-based approaches consider additional penalties or constraints during training to maintain the memory of old tasks. For example, Learning without Forgetting (LwF) exploits the idea of knowledge distillation [14] to force the model outputs on the current data to be similar to the old model's. Similarly, *Pseudo-Label* also

¹<https://github.com/marinacecon1/RCLP>




Figure 2. Benchmark proposed for the medical imaging field. The figure presents a New Instances & New Classes scenario, where some tasks introduce new classes while other tasks involve previously seen classes but with a shift in the input data distribution. Our proposed stream consists of a sequence of seven tasks, encompassing a total of nineteen classes across two domains.

uses the model trained on the previous task [13]. While visiting the data of a new task, the new task samples are passed through the old model, and the predictions on the old labels are added to the new ground truth targets.

Architecture-based approaches, as the name suggests, are methods that alter the original model’s architecture to preserve previous knowledge. Techniques within this category employ various strategies for modifying the architecture [9, 31, 37].

Moreover, in CL literature, some methods propose combining different families of strategies to obtain better and more robust performance. For example, Hou et al. [15] combine a distillation-based approach with Replay. Another famous example is Dark Experience Replay (DER) [2]; rather than the ground truth targets, this method stores the logits produced from the previous models. In this way, the knowledge is distilled from the previously optimized models while replaying old input samples.

2.2. CL in the Medical Domain

Previous works for CL in the medical domain have explored DIL and CIL scenarios within CL methods for classifying chest X-ray images. For the DIL scenario, Lenga et al. [26] conduct a study considering two tasks based on chest X-ray images; Srivastava et al. [42] address a similar setting through a replay-based approach. Singh et al. [40] consider a CIL scenario with three tasks encompassing 12 classes, gradually introducing new classes over time. Similarly, Akundi et al. [1] study a CIL scenario with five tasks, each with only one class, for a total of 5 classes considered. The primary challenge in the Class Incremental Learning (CIL) scenario within this context is its multi-label characteristic, leading to the label absence problem. Specifically, images in each task may contain multiple classes, but only the labels relevant to the current task are provided. This issue is also encountered in Object Detection [33, 39] and Semantic Segmentation [3, 7, 12], where it is known as “background shift”. It is particularly problematic for Replay-based methods, as a single sample may have different labels across tasks, leading to interference. Consequently, many chest X-ray studies prefer distillation-based approaches [1].

However, previous works found a significant degradation in the performance of these approaches when assuming that certain classes encountered at the start of the data stream rarely reappear in subsequent tasks [19]. These considerations bring additional challenges that previously proposed methods could not handle correctly.

3. Considered Scenario

3.1. Motivation

We consider the problem of flexible collaboration across medical institutions, enabling effective data pooling and model sharing. The nature of healthcare systems and processes leads to three challenges. First, medical data is highly sensitive and subject to stringent data protection regulations [11]. Sharing and processing medical data across multiple institutions is especially difficult as hospitals are independent organizations with distinct data controllers [8]. Informed consent can enable broader processing of medical data, but it is challenging to obtain for a large and representative group of patients [5].

Furthermore, medical diagnosis encompasses numerous disease labels, necessitating the curation and annotation of a comprehensive dataset across multiple pathologies, which can take several years [24]. It is crucial for models to evolve incrementally and incorporate information about new diseases as they become available [43] or surface, as was the case with COVID-19.

Moreover, the classification of rare pathologies often benefits from extending models to related yet different conditions [34]. Typically, access to the original data used for training the model is not feasible due to data protection regulations. Consequently, a scenario where the model is updated with only a small subset of historical data, like the CIL setting, becomes crucial.

Finally, domain shift is inevitable, e.g., due to differences in medical equipment, acquisition protocols, and evolving populations [20, 26, 41]. This is especially true for datasets spanning multiple sites, which may serve areas with different demographics or provide specialized care, e.g., to different age ranges. In general, medical processes

adopted at one hospital influence data acquisition and cause artifacts in the data [32].

Overall, both DIL and CIL scenarios are common in the medical domain. Clearly, they do not represent mutually exclusive settings. Hence, a situation in which both new domains and new classes appear is not only possible but also very plausible, and it needs to be addressed. For this reason, we consider a scenario termed **New Instances & New Classes**. NIC combines the challenges of both new class arrivals and domain shifts within a single framework, mirroring the complexities often encountered in realistic applications like medical imaging. This combination of different types of shifts between tasks makes this scenario more challenging and suitable for representing intricate settings, such as the medical domain. The concept of the NIC scenario was initially introduced and investigated by Lomonaco et al. [30]. It was examined within the context of a single-label classification problem, where it is assumed there are no intersections among tasks regarding labels or samples. To the best of our knowledge, this scenario was never analyzed in multi-label settings or the medical imaging setting.

3.2. Scenario Description

In this work, we model a NIC scenario in chest X-ray imaging across two sites. We consider the ChestX-ray14 dataset (NIH), compiled from patients of the NIH Clinical Center [47] and the CheXpert dataset (CXP), with radiographies from patients treated at the Stanford Hospital [17]. Both datasets contain information on 14 diseases, 7 of which are in common. We removed non-pathology labels from CXP, resulting in a stream of 19 classes. We divided CXP into three tasks and the NIH dataset into four tasks, ordering them sequentially to create the NIC stream of tasks represented in Figure 2. Each task models the availability of new data at one hospital. Between tasks i and $i + 1$, either new classes are introduced, or data from previously known classes is presented with a shift in the input data distribution.

In each task, all and only the images containing the associated pathologies in the corresponding dataset are present. Since several images contain multiple labels, the intersection among tasks is not null, presenting a realistic scenario where images with previously seen (but currently not labeled) diseases can appear in new tasks.

3.3. Challenges

The proposed benchmark considers a multi-label image classification problem where new classes and domains are added over time. We identify these three challenges in the current literature landscape of multi-label classification: (i) the task interference problem in replay approaches, (ii) the potential not exploited of the replay memory samples, (iii) the strong forgetting suffered by distillation methods if pre-

viously seen classes don't reappear in the current task.

Indeed, while replay-based approaches have shown remarkable performance in multi-class classification problems [49], they encounter a specific challenge known as *task interference* in multi-label continual learning (CIL) settings. This interference arises between the current task samples and the replayed samples, when the intersection between tasks is not null, as each sample only contains information about the labels of the relative task. The same problem is encountered when solving other CL problems in multi-label settings, such as Object Detection [39] and Semantic Segmentation [19], where it is referred to as "foreground shift" [50].

Moreover, we state that another problem of the Replay approach is that the replay memory is not fully exploited. Indeed, for each stored image, only the labels seen during its task are available, while all other labels are assumed to be unknown or absent. However, this hides the potential of the memory, which could contain much more information than what is actually stored.

Distillation-based methods have demonstrated suboptimal performance in single-label classification tasks [26]. Conversely, they find extensive application in other problems such as semantic segmentation [3] and object detection [39]. However, distillation methods demonstrate significant forgetting when previously encountered classes do not reappear in the current task, a scenario often unrealistic in practical applications. In other words, unlike replay methods, distillation-based methods used in multi-label settings perform better when there is a considerable intersection between tasks.

In the considered scenario, all these challenges are present. There is some intersection between tasks, causing the interference issue in Replay. However, the diseases of the first tasks are rare, hence they do not appear frequently in subsequent tasks, thereby challenging distillation methods. Lastly, since the considered setting is multi-label, the problem of the under-exploitation of the replay buffer is present.

4. Proposed approach: Replay Consolidation with Label Propagation

To tackle the challenges described in Sec. 3.3, we propose Replay Consolidation with Label Propagation. The advantages of our approach are threefold: (1) the integration of old knowledge on the new task samples and of new knowledge on the replay buffer samples, together with a Masking Loss, mitigate the issue of task interference of Replay; (2) the Replay memory is optimized since the targets provide information not only on the labels from their originating tasks but also from all preceding tasks up to the current one; (3) by replaying samples that contain old classes,

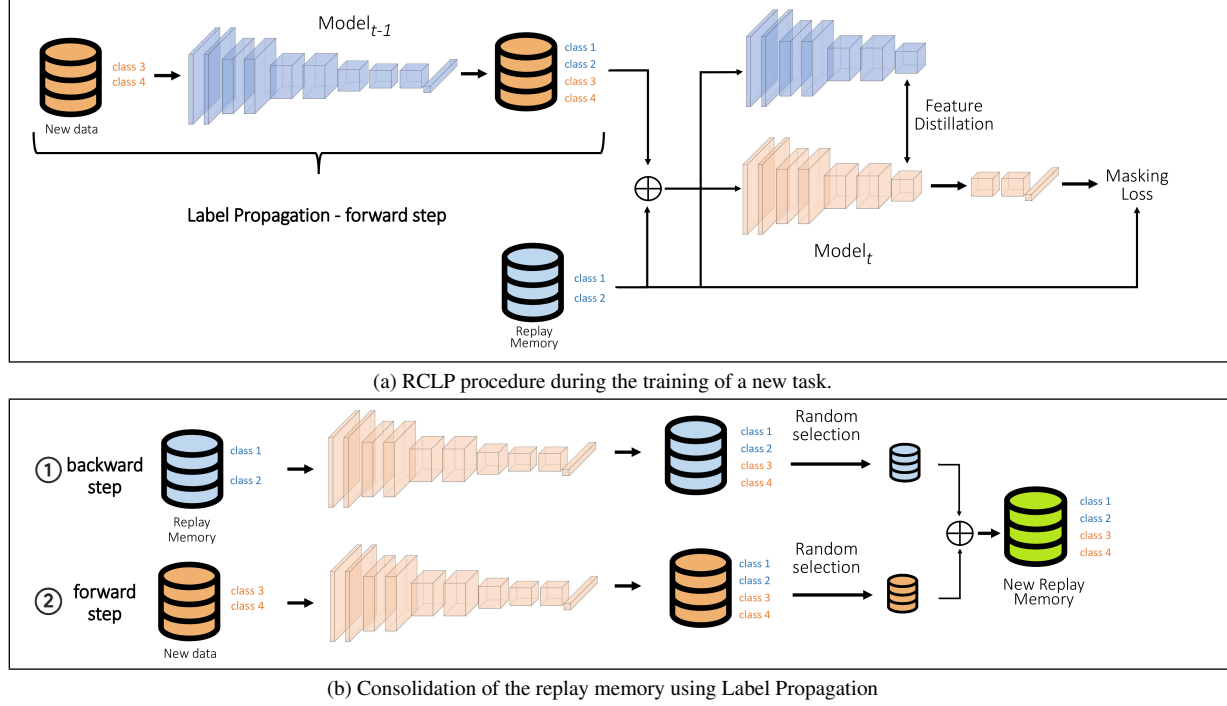


Figure 3. Representation of the RCLP method: (a) New task samples and those from the replay memory buffer are used for training. New task samples are first processed through the old model to generate pseudo-labels. The loss for replayed samples is masked to focus on old task labels and combined with classification and intermediate losses. (b) Memory buffer consolidation via label propagation. The backward step updates the buffer with new knowledge, while the forward step integrates old knowledge into new task samples. Some of these updated samples are then saved to the memory buffer after training.

the limitations of distillation are overcome. The comprehensive strategy is depicted in Fig. 3.

4.1. Label Propagation: Forward step

The first issue we address is that during the training of the model on new samples, these samples contain information only about the new labels and lack information on labels seen in previous tasks. Therefore, in this part, we aim to incorporate the information on the old labels into the new data, and subsequently in the replay buffer, via the pseudo-labeling technique.

This step can be defined formally as follows. Let X_t, Y_t be the data associated with task t . Let X_M, Y_M be the data associated with the Replay memory. Let f_{θ_t} be the model after being trained on task t . Let L_i be the set of labels associated with task i , $y_j^{L_i}$ the output produced by the model concerning the labels L_i , and $y_j^{L_i}$ the j -th output from the labels L_i . During the training of task t , given a sample $x \in X_t$, the forward step adjusts the ground truth associated with each $y \in Y_t$, relative to the labels L_1, \dots, L_{t-1} , to integrate the knowledge of the previously optimized model $f_{\theta_{t-1}}$. To achieve this goal, we utilize the predictions generated by the previous model, denoted as $\hat{y}^{L_1}, \dots, \hat{y}^{L_{t-1}}$. These predictions undergo a thresholding process where

each class label is determined. Subsequently, these predicted labels replace the outdated labels in the ground truth vector y .

$$\forall y \in Y_t \quad y_j^{L_i} = \begin{cases} 1, & \text{if } \hat{y}_j^{L_i} > h \\ 0, & \text{if } \hat{y}_j^{L_i} \leq h \end{cases} \quad \begin{matrix} 1 \leq j \leq |L_i| \\ 1 \leq i \leq t-1 \end{matrix} \quad (1)$$

After training on the new task t , a subset of the samples from task t is saved in the memory buffer M . The targets associated with these samples, relative to old tasks labels, are determined using the label propagation technique previously described. This approach ensures that, following the training on task t , the memory buffer contains samples with information relevant to all tasks up to and including task t .

4.2. Label Propagation: backward step

As described in section 3.3, the main problem of the replay approach is that the samples for each task contain only the labels of said task, as shown in Fig. 4a. This represents an underuse of the potential information of the data contained in memory, which could be able to express much more, and it's a partial cause of the interference issue. We partially solve this problem by performing the forward step

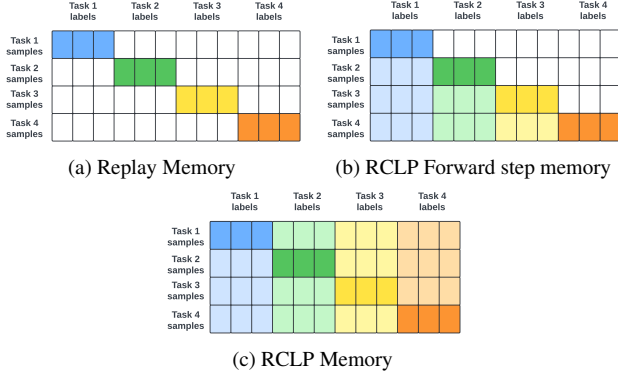


Figure 4. Representation of replay memories for different approaches. The vertical axis represents the samples of each task, while the horizontal axis is the labels associated with each task. (a) Each task of the Replay memory has information only on the labels seen during its iteration. (b) The forward step of label propagation, saves samples that are informative of all tasks up to the current one. (c) The backward step integrates the knowledge of the new labels in the replay buffer.

of Label Propagation when creating the new replay memory. In this way, samples from task i in the memory buffer are informative on all tasks up to task i , as shown in figure 4b.

While this approach reduces the interference issue, the samples in the memory buffer do not contain any information about the tasks that come after their respective task of origin. Therefore, we propose the backward step of Label Propagation to consolidate the new knowledge acquired after training on a new task in the old samples of the replay buffer, as shown in Fig. 4c.

The procedure can be described formally as:

$$\forall y \in Y_M \quad y_j^{L_t} = \begin{cases} 1, & \text{if } \hat{y}_j^{L_t} > h \\ 0, & \text{if } \hat{y}_j^{L_t} \leq h \end{cases} \quad 1 \leq j \leq |L_t| \quad (2)$$

In other words, after training the model on the new task t , the samples in the replay buffer, coming from previous tasks, are given as input to the model, and the pseudo-labels relative to the current task are computed and added to the ground truth targets.

4.3. Masking Loss

After performing the label propagation step, only the labels relative to the current task in the replayed samples are absent. Therefore, we propose a masking loss to overcome this issue. When the model revisits old samples, the masking loss ensures that the model’s outputs are only influenced by the labels of the old tasks. This effectively ‘masks’ the influence of the old samples on the new labels, preventing

task interference.

$$L_M(x, y, \hat{y}) = \begin{cases} L_{BCE}(y, \hat{y}), & x \in X_t \\ \sum_{i=1}^{t-1} L_{BCE}(y^{L_i}, \hat{y}^{L_i}), & x \in X_M \end{cases} \quad (3)$$

4.4. Feature Distillation

To further improve the model’s performance, we consider the feature distillation technique on the intermediate representations of the model [28]. Let our model f be composed as $g_\phi(h_\omega(x))$, where given x , h produces the features, and the classifier g is trained on such features. The main idea of feature distillation is that the features produced by the new model on the new and old data should not deviate too far from the features produced by the previous model on said data. This allows the model to maintain a more consolidated memory mitigating the effects of catastrophic forgetting. Formally, given an input sample x :

$$L_{FD} = ||g_{\phi_t}(h_{\omega_t}(x)) - g_{\phi_{t-1}}(h_{\omega_{t-1}}(x))|| \quad (4)$$

The input could either come from the new task or from the replay buffer. Therefore the final loss is built in the following way:

$$L(x, y, \hat{y}) = L_M(x, y, \hat{y}) + \gamma \cdot L_{FD}(x) \quad (5)$$

4.5. Method Analysis

In conclusion, RCLP optimally leverages the strengths of both distillation and rehearsal methods. To address the problem of task interference, we split our RCLP approach into multiple components. We propose the Label Propagation technique, which can be divided into two steps. The **forward step**, like Pseudo-Label, enables new samples to leverage the knowledge from a previously optimized model to acquire the old labels. However, unlike Pseudo-Label, our approach applies this step not only during training but also to consolidate information within the replay memory buffer. Instead, the **backward step** allows the injection of the new knowledge (acquired after the training on the new task) into the samples of replay memory. In this way, a relevant advantage of Label Propagation is represented by the full exploitation of Replay memory by utilizing samples that contain information on all seen tasks (as depicted by Fig. 4). Moreover, we introduce a *Masking Loss* to further mitigate the task interference issue by considering only the old labels in the replay samples, hence avoiding the calculation of loss on the labels of the current task. In addition, to enhance final performance and reduce further forgetting, a feature distillation loss is implemented.

Moreover, unlike other methods [7] that update replayed samples with the most recent model, which can lead to unreliable predictions due to forgetting, our method updates replayed sample labels only once, immediately after training on the current task, ensuring more robust pseudo-labels.

Additionally, the computational cost of RCLP is similar to that of methods like DER and LwF Replay [2, 15]. It requires storing a subset of past samples (as in Replay) and weights of previous models (as in distillation). Label propagation does not increase memory use, as pseudo-labels replace original labels, and the masking and intermediate loss procedures do not add extra overhead. While replay increases computation time compared to Fine-Tuning, the label propagation steps are minimal. Each minibatch passes through both old and new models, similar to LwF, without significantly affecting computation time. Processing the memory buffer after each task is also efficient, as its small size makes it comparable to computing a validation loss.

5. Experimental Setting

We test the following methods against the benchmark: Joint training, Fine tuning, Replay, LwF, Pseudo-Label, LwF-Replay, DER. Following the literature on training classification models on chest X-ray images [38], we selected a 121-layer DenseNet [16] architecture pre-trained on ImageNet [6] for our study. We employed the binary cross-entropy loss function and optimized the models using the Adam [22] optimizer with a learning rate set at 0.0005. For both datasets, we used an 80-10-10 split between training, validation, and test sets. Consistent with findings from previous works [48] indicating that using a single image per patient does not significantly impact performance, we included only one image per patient, prioritizing frontal images. All images were resized to 256×256 and normalized via the mean and standard deviation of the ImageNet dataset. As an upper bound for our methods, we considered Joint Training, which consists in simultaneously training on all tasks. In this case, we reduced the learning rate by a factor of 2 if the validation loss did not improve over three epochs and stopped training if the validation loss did not improve over 10 epochs. For all other methods besides Joint training, we trained on each task for 10 epochs. As a lower bound, we evaluated the Fine-Tuning approach, where each task is trained sequentially, with no additional technique to avoid catastrophic forgetting. For the Replay approach, we consider a mix ratio of 50% and a memory size corresponding to 3% of the original dataset size. The samples to store in the memory buffer and to propose to the model in future tasks are chosen uniformly at random. In LwF [29], the loss is computed as $L = L_1 + \tau L_2$, where L_1 is the loss for the current task, while L_2 is the distillation loss; we set $\tau = 2$, following the literature [29]. For Pseudo-Label, we define a different threshold for each class, choosing the one found to maximize the F1 score on the validation sets relative to the origin task of each label. When no new labels are introduced between task i and task $i + 1$, the thresholds remain fixed at the values that were optimal in the previous validation sets. For the LwF-Replay approach, we used the same

hyperparameters and sampling strategy of Replay and LwF. Similarly, for the DER approach, the same hyperparameters of Replay are used. Finally, for our approach, we adopt the same hyperparameters and sampling strategy as in Replay and employ the same technique for computing thresholds as in Pseudo-Label. Moreover, for computing the intermediate distillation loss, we utilized the output of the 12th block out of the 16 blocks in the DenseNet architecture. Lastly, the parameter γ of the feature distillation loss in Eq. 4.4 is set to 1.

Evaluation metrics We use the macro F1 score as the primary metric for evaluating model performance, as it is common in multi-label classification [10, 46]. The macro F1 averages the F1 score for each label, balancing precision and recall, making it ideal for imbalanced datasets. Its focus on positive cases, like detecting pathologies in medical images, adds relevance in this domain. We also report the AUC ROC, a popular metric for medical image classification [26, 38, 42], though it can be misleading on imbalanced datasets due to its limitations in reflecting poor model performance [18].

For the evaluation over a stream of tasks, we consider the Average F1 over all the pathologies of all tasks up to i , as typical in CL papers [1, 26]. Similarly, we measure the percentual forgetting, proposed in previous papers in CL literature [4, 21], hence the percentage difference between the final performance of the model on each task and the performance just after training on said task. Lastly, we consider the Relative gap between each method’s final performance and the upper baseline’s performance, i.e., Joint Training. The formulas for computing these two metrics are reported in the Appendix.

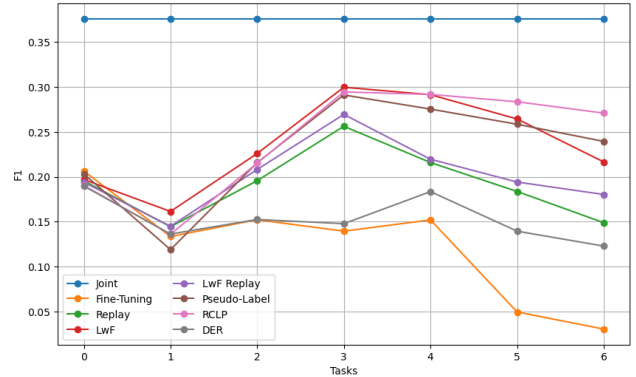


Figure 5. F1 score during the task stream for each method.

6. Results

In the next part, we discuss the outcomes produced by each method when applied to the proposed challenging benchmark. Specifically, we will elucidate the achieved re-

Table 1. In **bold** is reported the best value for each metric among the compared CL approaches. Average F1, Forgetting, and Relative gap metrics are as defined in Sec. 5.

Strategy	Metrics			
	Avg. F1 \uparrow	Avg. AUC \uparrow	Forgetting F1 \downarrow	Relative gap \downarrow
Joint Training	0.38	0.79	-	-
Fine-Tuning	0.02	0.55	100%	95%
Replay	0.15	0.65	59%	60%
LwF	0.22	0.68	41%	42%
LwF Replay	0.18	0.68	48%	52%
DER	0.12	0.58	73%	68%
Pseudo-Label	0.24	0.69	21%	37%
RCLP (ours)	0.27	0.69	2.4%	29%

sults concerning the F1 performance in Sec. 6.1. Subsequently, we delve deeper by examining the forgetting behavior exhibited by each method in Sec. 6.2.

6.1. Performance results

The results are reported in Tab. 1, while the average F1 metric over time for each method is shown in Fig. 5. The reported results are averaged over multiple repetitions to reduce the effect of randomness on the experimental results. Despite the issue of task interference, Replay has an average F1 score of 0.15, still much higher than the lower bound Fine-Tuning that achieves an F1 of 0.02. This discrepancy in performance between the two methods can be observed from the relative gap as well: 95% for Fine-Tuning and 60% for Replay. LwF outperforms Replay, as expected from the literature [39], exhibiting a significant improvement, achieving an average F1 score of 0.22. Pseudo-Label performs similarly to LwF, with a marginally higher F1 score of 0.24. Interestingly, the combination of LwF and Replay, LwF-Replay, isn't able to produce satisfactory results, obtaining an F1 value (0.18) lower than that of the LwF and Pseudo-Label methods taken singularly. This is expected due to the interference caused by Replay, which is not overcome by combining the method with LwF. For the same reasons, DER performs very poorly as well, achieving a final F1 of 0.12, lower than all methods besides Fine-Tuning. Indeed, as discussed, a series of challenges need to be addressed to exploit the advantages of replay-based methods together with distillation-based methods in the multi-label setting. Lastly, our approach, Replay Consolidation with Label Propagation, achieves an F1 score of 0.27 and a relative gap of 29%, exhibiting a notable improvement with respect to the other methods, as shown in Tab. 6 and Fig. 5).

6.2. Forgetting Analysis

Another interesting analysis concerns the forgetting metric (shown in Tab. 1). The first general insight is that the

forgetting values are notably high across all methods besides RCLP, ranging from 73% for DER to 21% for Pseudo-Label. When examining our proposed approach, RCLP, we note a remarkably low forgetting rate of 2.4%, with the final performance clearly outperforming all other methods. This is also supported by the results in the Appendix.

7. Conclusions and Future Work

Our work introduces a novel benchmark tailored for evaluating CL methods in the domain of multi-label medical image classification. This benchmark includes different medical imaging datasets, pathologies, and imaging conditions. In particular, we designed the benchmark by considering realistic conditions in the medical settings and combining new classes and domains in the task stream.

Because of the challenges encountered in the scenario, we proposed a novel method called Replay Consolidation with Label Propagation (RCLP). Our method outperforms existing approaches in the field of multi-label image classification within the medical domain, demonstrating significantly superior performance with minimal forgetting.

While our approach has demonstrated high effectiveness in handling multi-label classification problems within the medical setting, its applicability to other settings still needs to be explored. In future work, we plan to evaluate the effectiveness of RCLP in other scenarios, such as object detection and semantic segmentation problems. Additionally, we intend to evaluate variations of the proposed method, such as exploring alternative approaches for computing the thresholds used in Pseudo-Labeling, as well as different strategies for selecting the samples to be stored in the replay buffer.

References

- [1] Prathyusha Akundi and Jayanthi Sivaswamy. Incremental learning for a flexible cad system design. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–4. IEEE, 2022. 3, 7

- [2] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline, 2020. 3, 7
- [3] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulò, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation, 2020. 3, 4
- [4] Davide Dalle Pezze, Denis Deronjic, Chiara Masiero, Diego Tosato, Alessandro Beghi, and Gian Antonio Susto. A multi-label continual learning framework to scale deep learning approaches for packaging equipment monitoring. *Engineering Applications of Artificial Intelligence*, 124:106610, 2023. 7
- [5] Fida K Dankar, Marton Gergely, and Samar K Dankar. Informed consent in biomedical research. *Computational and structural biotechnology journal*, 17:463–474, 2019. 3
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 7
- [7] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Tackling catastrophic forgetting and background shift in continual semantic segmentation, 2021. 3, 6
- [8] Edward S Dove. The eu general data protection regulation: implications for international scientific research in the digital era. *Journal of Law, Medicine & Ethics*, 46(4):1013–1030, 2018. 3
- [9] 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
- [10] Weifeng Ge, Sibe Yang, and Yizhou Yu. Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1277–1286, 2018. 7
- [11] Lawrence O Gostin. National health information privacy: regulations under the health insurance portability and accountability act. *Jama*, 285(23):3015–3021, 2001. 3
- [12] Dipam Goswami, René Schuster, Joost van de Weijer, and Didier Stricker. Attribution-aware weight transfer: A warm-start initialization for class-incremental semantic segmentation, 2022. 3
- [13] 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
- [14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [15] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Lifelong learning via progressive distillation and retrospection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 3, 7
- [16] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks, 2018. 7
- [17] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019. 4
- [18] László A Jeni, Jeffrey F Cohn, and Fernando De La Torre. Facing imbalanced data—recommendations for the use of performance metrics. In *2013 Humaine association conference on affective computing and intelligent interaction*, pages 245–251. IEEE, 2013. 7
- [19] Tobias Kalb, Masoud Roschani, Miriam Ruf, and Jürgen Beyerer. Continual learning for class- and domain-incremental semantic segmentation. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 1345–1351, 2021. 1, 3, 4
- [20] Neerav Karani, Krishna Chaitanya, Christian Baumgartner, and Ender Konukoglu. A lifelong learning approach to brain mr segmentation across scanners and protocols. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 476–484. Springer, 2018. 3
- [21] Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. Imbalanced continual learning with partitioning reservoir sampling. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII* 16, pages 411–428. Springer, 2020. 7
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 7
- [23] 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
- [24] Pratyush Kumar and Muktabh Mayank Srivastava. Example mining for incremental learning in medical imaging. In *2018 IEEE symposium series on computational intelligence (SSCI)*, pages 48–51. IEEE, 2018. 3
- [25] Pratibha Kumari, Joohi Chauhan, Afshin Bozorgpour, Reza Azad, and Dorit Merhof. Continual learning in medical imaging analysis: A comprehensive review of recent advancements and future prospects. *arXiv preprint arXiv:2312.17004*, 2023. 1
- [26] Matthias Lenga, Heinrich Schulz, and Axel Saalbach. Continual learning for domain adaptation in chest x-ray classification. In *Medical Imaging with Deep Learning*, pages 413–423. PMLR, 2020. 1, 3, 4, 7
- [27] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges, 2019. 2
- [28] Songze Li, Tonghua Su, Xuyao Zhang, and Zhongjie Wang. Continual learning with knowledge distillation: A survey. *Authorea Preprints*, 2024. 6
- [29] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 7

- [30] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition, 2017. 1, 4
- [31] 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
- [32] Sinan Özgün, Anne-Marie Rickmann, Abhijit Guha Roy, and Christian Wachinger. Importance driven continual learning for segmentation across domains. In *Machine Learning in Medical Imaging: 11th International Workshop, MLMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 11*, pages 423–433. Springer, 2020. 4
- [33] Francesco Pasti, Marina Ceccon, Davide Dalle Pezze, Francesco Paissan, Elisabetta Farella, Gian Antonio Susto, and Nicola Bellotto. Latent distillation for continual object detection at the edge. In *European conference on computer vision*, 2024. 3
- [34] Gwenolé Quéllec, Mathieu Lamard, Pierre-Henri Conze, Pascale Massin, and Béatrice Cochener. Automatic detection of rare pathologies in fundus photographs using few-shot learning. *Medical image analysis*, 61:101660, 2020. 3
- [35] ANTHONY ROBINS. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. 2
- [36] 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
- [37] 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
- [38] Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, Irene Chen, and Marzyeh Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. pages 232–243, 11 2020. 7
- [39] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, pages 3400–3409, 2017. 3, 4, 8
- [40] Amritpal Singh, Mustafa Burak Gurbuz, Shiva Souhith Gantha, and Prahlad Jasti. Class-incremental continual learning for general purpose healthcare models. *arXiv preprint arXiv:2311.04301*, 2023. 1, 3
- [41] Mehreen Sirshar, Taimur Hassan, Muhammad Usman Akram, and Shoab Ahmed Khan. An incremental learning approach to automatically recognize pulmonary diseases from the multi-vendor chest radiographs. *Computers in Biology and Medicine*, 134:104435, 2021. 3
- [42] Shikhar Srivastava, Mohammad Yaqub, Karthik Nandakumar, Zongyuan Ge, and Dwarikanath Mahapatra. Continual domain incremental learning for chest x-ray classification in low-resource clinical settings. In *MICCAI Workshop on Domain Adaptation and Representation Transfer*, pages 226–238. Springer, 2021. 1, 3, 7
- [43] Le Sun, Mingyang Zhang, Benyou Wang, and Prayag Tiwari. Few-shot class-incremental learning for medical time series classification. *IEEE Journal of Biomedical and Health Informatics*, 2023. 3
- [44] Qazi Mazhar ul Haq, Shanq-Jang Ruan, Muhammad Amirul Haq, Said Karam, Jeng Lun Shieh, Peter Chondro, and De-Qin Gao. An incremental learning of yolov3 without catastrophic forgetting for smart city applications. *IEEE Consumer Electronics Magazine*, 11(5):56–63, 2021. 1
- [45] Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. *CoRR*, abs/1904.07734, 2019. 2
- [46] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2285–2294, 2016. 7
- [47] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2097–2106, 2017. 4
- [48] Nina Weng, Siavash Bigdeli, Eike Petersen, and Aasa Fera-gen. Are sex-based physiological differences the cause of gender bias for chest x-ray diagnosis? In *Workshop on Clinical Image-Based Procedures*, pages 142–152. Springer, 2023. 7
- [49] Qihan Yang, Fan Feng, and Rosa Chan. A benchmark and empirical analysis for replay strategies in continual learning, 2022. 4
- [50] Liu Yuyang, Cong Yang, Goswami Dipam, Liu Xialei, and Joost van de Weijer. Augmented box replay: Overcoming foreground shift for incremental object detection, 2023. 4
- [51] Haoran Zhang, Natalie Dullerud, Karsten Roth, Lauren Oakden-Rayner, Stephen Robert Pfohl, and Marzyeh Ghassemi. Improving the fairness of chest x-ray classifiers, 2022. 1