



Continuous fake media detection: adapting deepfake detectors to new generative techniques

Francesco Tassone^a, Luca Maiano^{a,b}, Irene Amerini^a

^aSapienza University of Rome, Via Ariosto 25, 00185 Rome, Italy

^bUbiquitous srl, Via Ariosto 25, 00185 Rome, Italy

ABSTRACT

Generative techniques continue to evolve at an impressively high rate, driven by the hype about these technologies. This rapid advancement severely limits the application of deepfake detectors, which, despite numerous efforts by the scientific community, struggle to achieve sufficiently robust performance against the ever-changing content. To address these limitations, in this paper, we propose an analysis of two continuous learning techniques on a *Short* and a *Long* sequence of fake media. Both sequences include a complex and heterogeneous range of deepfakes (generated images and videos) from GANs, computer graphics techniques, and unknown sources. Our experiments show that continual learning could be important in mitigating the need for generalizability. In fact, we show that, although with some limitations, continual learning methods help to maintain good performance across the entire training sequence. For these techniques to work in a sufficiently robust way, however, it is necessary that the tasks in the sequence share similarities. In fact, according to our experiments, the order and similarity of the tasks can affect the performance of the models over time. To address this problem, we show that it is possible to group tasks based on their similarity. This small measure allows for a significant improvement even in longer sequences. This result suggests that continual techniques can be combined with the most promising detection methods, allowing them to catch up with the latest generative techniques. In addition to this, we propose an overview of how this learning approach can be integrated into a deepfake detection pipeline for continuous integration and continuous deployment (CI/CD). This allows you to keep track of different funds, such as social networks, new generative tools, or third-party datasets, and through the integration of continuous learning, allows constant maintenance of the detectors.

© 2024 Elsevier Ltd. All rights reserved.

1. Introduction

Generative AI tools like Midjourney¹, ChatGPT², or the more recent Sora³ are completely revolutionizing the way media content is created, leading to the mass adoption of tools that were unimaginable until recently. However, this progress makes the threat of new-generation disinformation or defamation campaigns increasingly concrete. Unfortunately, forensic

tools for detecting this content are not advancing at the same rate. Although it is possible to train highly accurate detectors, these methodologies still poorly generalize to new generative methods due to data drift (Paley et al. (2022)). Detectors perform well on the generative techniques they are trained on but commonly fail when exposed to content generated with a new generative model.

Because of these limitations, the practical application of automatic detectors has been almost nil. When someone wants to deploy these tools in commercial or mass verification systems, they must face numerous challenges that go far beyond the need to generalize from a few known benchmarks (Rossler et al. (2019); Corvi et al. (2023); Li et al. (2023); Zi et al. (2020)). Prominent among these is the need to continuously

e-mail: tassone.1814263@studenti.uniroma1.it (Francesco Tassone), maiano@diag.uniroma1.it (Luca Maiano), amerini@diag.uniroma1.it (Irene Amerini)

¹<https://www.midjourney.com/>

²<https://openai.com/chatgpt>

³<https://openai.com/sora>

train these models on new generative techniques through continuous learning methods. To address this continuous change, we analyze continuous learning techniques to evaluate their limitations when applied to mitigate this problem. Continual learning, also known as lifelong learning or incremental learning, is a constant learning approach. Unlike transfer learning techniques, where a model trained on one task is retrained on a new task to improve its performance on the latter, continual learning involves maintaining good model performance on a set of evolving tasks as these become available without incurring in *catastrophic forgetting* (French (1999)). These requirements fit well with the problem of learning to recognize content generated by new techniques and the need to continually readjust a model with respect to data drift produced by the shift in the distribution of data observed in inference versus those seen in training.

This work aims to investigate the effectiveness of two continuous learning techniques with the intention of integrating them into a real, end-to-end deepfake detection system that allows for continuous integration and continuous delivery/deployment (CI/CD). Our goal is therefore to propose a simple yet effective design for a Machine Learning Model Operations (MLOps Paleyes et al. (2022); Semola et al. (2022)) pipeline that would enable the end-to-end development of continuously trained and monitored intelligent detectors with a minimal set of components.

Contributions

Therefore, the main contributions of this preliminary study are summarized below.

- *Analysis of continual learning methods.* We study the effectiveness of two continuous learning methods, *Knowledge Distillation KD* – (Hinton et al. (2015)) and *Elastic Weight Consolidation (EWC* – Kirkpatrick et al. (2017)), and show their superiority to transfer learning when continuous training is needed.
- *Sequence.* We study how the order of arrival of the tasks can affect the performance of the model. In particular, we show how task similarity plays an important role in maintaining optimal performance.
- *Multi-task continual training.* We show how aggregating tasks based on their similarity can significantly improve the overall performance over the entire sequence. This result is particularly important and helps us to better outline the possible developments of these techniques for deepfake recognition.
- *CI/CD pipeline for deepfake detection.* We propose an overview of an end-to-end system for continuous integration and continuous delivery/deployment for a deepfake detection application.

The rest of this paper is organized as follows. Section 2 offers an overview of the state of the art. In Section 3, we introduce the methodology used in this study. In Section 4, we present the experiments we conducted. Finally, in Section 5, we draw

the final considerations and illustrate the future developments of this work.

2. Background and related works

2.1. Machine Learning Operations (MLOps)

MLOps is a set of scientific principles, tools, and techniques for designing and constructing complex computer systems that combine the scientific ideas, instruments, and methods of traditional Software Engineering (DevOps, Leite et al. (2019)) with Machine Learning. It covers every phase, including gathering data, creating the model, and integrating it into the software production system. This is an actual engineering field that deserves detailed study; however, for this reason, in the following paragraphs, we will provide a very brief introduction.

MLOps results from the realization that the quality, transparency, and agility of the entire intelligent software are reduced when the ML model development process is isolated from the ML operations process that produces it⁴. A typical machine learning-based software includes four key elements: (1) the data, (2) the machine learning model, (3) the code, and (4) performance assessment, and each of these elements contributes to creating four main phases. The *data engineering* phase involves the study of gathering and preparing data. *machine learning model engineering* is a process that begins with model evaluation, training, and serving. *Code engineering* works on best practices to integrate the model into the final product. Finally, *performance monitoring* deals with performance assessment, reproducibility, and monitoring in production to find deviations.

The quality and applicability of machine learning models in production are affected by three main issues. First, the quality of input data (quantity, completeness, and semantics) impacts the performance of machine learning models. The second is model decay, which is the gradual deterioration of the performance of the machine learning models. In reality, real-world data is not stationary and is often absent from model training. In this paper, we focus on this specific aspect. Third, according to quality metrics (location). Models pre-trained on known training data may not perform correctly when applied to new inputs. This aspect is essential for deepfakes, where new tools and generative techniques can undermine the quality of the detectors. Readers more interested in this topic can learn more about the challenges and commercially available MLOps tools in Hewage and Meedeniya and Kreuzberger et al. (2023).

2.2. Deepfake detection

Despite the difficulties in keeping up with the advancement of generative techniques, numerous studies have been proposed on detecting generated images and videos (Amerini et al. (2021); Verdoliva (2020)). In this section, we provide an overview of the most recent advances. However, it is essential to note that many can be combined with the continuous learning techniques analyzed in this paper. In fact, this type of learning

⁴<https://ml-ops.org/>

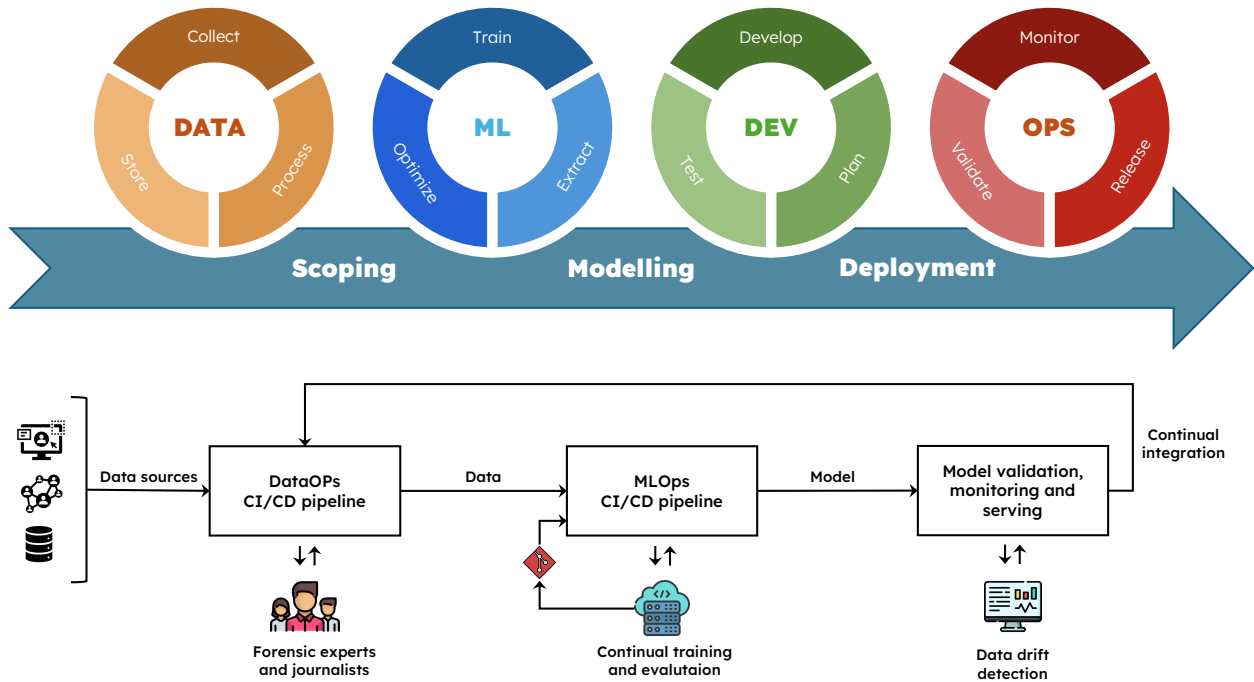


Fig. 1: Proposed CI/CD pipeline for deepfake detection. The data coming from different sources like new generative tools, social media, or existing databases are analyzed by forensic experts to ensure the system’s continual retraining. Next, these data are used for continual learning and monitoring. The data drift distribution module raises an alert whenever it detects new input data distributions. The continual learning methods analyzed in this paper are part of the *MLOps CI/CD pipeline* block in the figure.

approach could be used not as an alternative to these methods but to make these techniques maintainable over time.

First of all, several studies have found that there are some key ingredients for more robust detection (Wang et al. (2020); Gragnaniello et al. (2021)). Among them, image compression and resizing can severely mitigate model performance (Papa et al. (2023)). Therefore, to cope with these problems, it is usually recommended to avoid resizes, as they entail image re-sampling and interpolation, which may erase the subtle high-frequency traces left by the generation process and train models with different forms of augmentation. Moreover, working on local patches also appears to be important (Chai et al. (2020)), as well as analyzing both local and global features (Ju et al. (2022)). Ciamarra et al. (2024) examine the effects that deepfake production may have had on the properties of the entire scene at the time of acquisition. Specifically, the deepfake generation process may alter all of these intrinsic relations. When an image (video) is taken, the scene’s overall geometry (such as surfaces) and the acquisition method (such as illumination) establish a clear environment that is directly represented by the image pixel values. A descriptor that may be used to train a CNN for deepfake detection can be obtained by analyzing the features of the surfaces shown in the image.

A recent study from Aghasanli et al. (2023) showed that foundational models like ViT can effectively distinguish between authentic and counterfeit images, even when interpretability through prototypes is important. Additionally, the study demonstrated that classifiers with fine-tuned features consistently outperform those utilizing pre-trained weights when applied to cross-dataset domains. Another study by Le and

Woo (2023) considers quality factors to train robust detectors. The authors used an intra-model collaborative learning method to minimize the geometrical differences of images in various qualities at different intermediate layers. This idea, combined with an adversarial weight perturbation module, can be used to improve the robustness of the model against input image compression. Similarly, several studies (Corvi et al., 2023; Sha et al., 2023; Ojha et al., 2023) use CLIP as a feature extractor. Cocchi et al. (2023) propose an analysis of the robustness of this model to image augmentations, transformations, and other pre-processing steps, while Amoroso et al. (2023) evaluate the effects of different image feature extractors, presenting results on CLIP and OpenCLIP. Many recent studies also focus on combining different modes, such as audio and video (Raza and Malik (2023); Zhang et al. (2024)) or video and depth (Mariano et al. (2022)), as well as open-set recognition (Wang et al. (2023)).

Other studies focus on reconstructing fake artifacts introduced by generative models by considering second-order statistics in the spatial and frequency domains (Wang et al. (2020)). Corvi et al. (2023) showed that similar to GANs, diffusion models also give rise to visible artifacts in the Fourier domain and exhibit anomalous regular patterns in autocorrelation. In fact, synthetic and real images exhibit significant differences in the mid-high frequency signal content, observable in their radial and angular spectral power distributions.

Among the various detection strategies, watermarks have also been proposed (Fei et al. (2023); Wu et al. (2020)). Through the addition of special information within the image being generated, these watermarks can be used to verify the

generative model used to create content. For example, Zhao et al. (2023) proposed an encoder-decoder network to embed watermarks as anti-deepfake labels into the facial identity features. The injected label is entangled with the facial identity feature, so it will be sensitive to face swap translations and robust to conventional image modifications like resizing and compression. Fernandez et al. (2022) examined the robustness of watermarks when the image is subjected to transformations. A similar study has also been on detection methods specifically designed for facial manipulation. However, these solutions are limited in that they require a model to integrate the watermark into the content during the generation phase (Lu and Ebrahimi, 2024).

Our method, instead, falls into a different line of studies targeted at designing continual learning methods for deepfake detection. This learning approach has only been partially explored for this task. The study by Marra et al. (2019) was one of the first to propose a multi-task incremental learning method for GAN-generated images based on iCaRL (Rebuffi et al. (2017)). A similar approach was applied to videos in Khan and Dai (2021), while Pan et al. (2023) have recently proposed to learn semantically consistent representations across domains based on supervised contrastive learning and a carefully designed replay set. In contrast, Kim et al. (2021) combined a knowledge distillation method with a representation learning loss. However, these methods have been tested on datasets not explicitly designed for this type of learning approach, which very often consists of a limited number of generative techniques. To overcome this limitation, Li et al. (2023) recently introduced a new collection of deepfakes from known and unknown generative models. The proposed CDDDB dataset includes multiple evaluations on detecting an easy, hard, and long sequence of deepfakes with appropriate measures. For these reasons, in this paper, we focus our studies on this dataset by measuring the performance of the continual learning methods examined in this paper on this dataset. A recent study from Epstein et al. (2023) takes a similar direction by applying online learning to several generative tools available today.

Different from other studies, we analyze this specific learning method to integrate it into a CI/CD system appropriately designed for this task in the future.

3. Method

The practical application of deepfake detectors has been severely limited by the need to develop robust detectors with respect to new generative techniques. To overcome this problem, we propose an analysis of two continual learning strategies and an overview of a simple CI/CD pipeline that can complement this preliminary study. This pipeline (depicted in Figure 1) has the advantage of being constantly updated with respect to the latest detectors, which can be adapted for continuous learning as described below.

3.1. Learning strategies

To evaluate the effectiveness of continual learning in the deepfake recognition task, we examine two learning methods.

These techniques were chosen based on their demonstrated effectiveness in other tasks (De Lange et al. (2022)).

Knowledge Distillation (KD)

Distillation techniques were introduced by Hinton et al. (2015) in order to transfer knowledge from a neural network \mathcal{T} (the *teacher*) to a neural network \mathcal{S} (the *student*). The key idea behind knowledge distillation is that soft probabilities predicted by a network of trained "teachers" contain much more information about a data point than a simple class label. For example, if multiple classes are assigned high probabilities for an image, this could mean that the image must be close to a decision boundary between those classes. Forcing a student to mimic these probabilities should then cause the student network to absorb some of this knowledge that the teacher discovered above and beyond the information in training labels alone. To implement the KD strategy, we modify the cross-entropy loss \mathcal{L}_S by adding a regularization term (\mathcal{L}_D) as follows.

$$\mathcal{L}_{KD}(\theta) = \alpha \mathcal{L}_D + \beta \mathcal{L}_S \quad (1)$$

Here, α and β are scalar coefficients that control the balancing between the current and past tasks. The distillation loss \mathcal{L}_D uses the output of the teacher model to facilitate knowledge transfer from previous tasks to the new student model. As the teacher is non-trainable, its predictions are solely based on prior knowledge, producing soft labels for the training process of the student model. This term is computed as:

$$\mathcal{L}_D = \sum_{x_i \in X} \sigma(\mathcal{T}(x_i, \hat{y}_i), \tau) \log \sigma(\mathcal{S}(x_i, y_i), \tau) \quad (2)$$

where \hat{y} represents the output of \mathcal{T} , σ denotes the softmax function with temperature τ and \mathcal{S} represent the student network.

Elastic Weight Consolidation (EWC)

EWC remembers old tasks by selectively slowing down learning on weights that are important for these tasks. As shown in Kirkpatrick et al. (2017), learning from a task A to a task B , there exist many configurations of θ leading to the same performance. In fact, the over-parametrization of the model makes it more likely the existence of a solution θ_B^* for task B that is close to task A . Therefore, previous tasks' performances are kept by constraining, with a quadratic penalty, the parameters to stay in a region centered in θ_A^* of low error for task A . Formally, the function \mathcal{L} that we minimize in EWC is:

$$\mathcal{L}_{EWC}(\theta) = \mathcal{L}_B(\theta) + \sum_i \frac{\lambda}{2} F_i(\theta_i - \theta_{A,i}^*)^2 \quad (3)$$

where F is the Fisher information matrix, λ sets how important the old task is compared to the new one and i labels each parameter. When moving to a third task (i.e., task C), EWC will try to keep the network parameters close to the learned parameters of both task A and B . This can be enforced either with two separate penalties, or as one by noting that the sum of two quadratic penalties is itself a quadratic penalty.

3.2. Training procedure

The training procedure for both learning methods described in the previous section is summarized in Algorithm 1. Given a stream of AI-generated contents $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_n\}, n \geq 1$, at each training iteration t we train a model $g_t(x_t, \theta_t)$ on the actual $\mathcal{D}_t \in \mathcal{D}$. After the first training iteration on the first available batch of samples $\mathcal{D}_1 \in \mathcal{D}$, the model is trained on all next batches $\{\mathcal{D}_2, \dots, \mathcal{D}_n\} \in \mathcal{D}$ for all $n > 1$. Differently from transfer learning, however, the model is forced to minimize the loss function for both new and old examples without requiring training over the past data samples, therefore learning an optimal set of parameters θ_t for all observed inputs $(x_i, y_i), i \in \{0, \dots, t\}$.

Algorithm 1 Training procedure for continual learning methods.

Require: $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_n\}, n \geq 1$
 $\mathcal{D}_t = \mathcal{D}.pop()$
 $g_t(x_t, \theta_t) \leftarrow train(\mathcal{D}_t)$
while $!\mathcal{D}.isEmpty()$ **do**
 $\mathcal{D}_t = \mathcal{D}.pop()$
if strategy == KD **then**
 $\mathcal{T} \leftarrow g_t(x_t, \theta_t).copy()$
 $S \leftarrow train(\mathcal{T}, \mathcal{D}_t)$ \triangleright Train using the \mathcal{L}_{KD} loss in Eq. 1.
 $g_t(x_t, \theta_t) \leftarrow S$
else if strategy == EWC **then**
 $g_t(x_t, \theta_t) \leftarrow train(\mathcal{D}_t)$ \triangleright Train using the \mathcal{L}_{EWC} loss in Eq. 3.
end if
end while

3.3. Lightweight CI/CD for deepfake detection

We conclude this section by providing an overview of how the proposed methodology could be integrated into a continuous integration and continuous delivery pipeline. The Figure 1 shows the complete pipeline. The upper part of the figure shows all the phases of an MLOps CI/CD system. The lower part shows the complete pipeline into which the continuous learning methods analyzed in this paper can be integrated.

The pipeline consists of three main modules. In the first **module**, data from generative tools, social media or databases are organized and analyzed if necessary by forensic experts or journalists. This module enables the preparation of model training data. In the next module, the model continuously learns from incoming data. This module also allows you to keep several copies of the model in case you need to restore a previous version. Finally, in the last part, a continuous delivery and monitoring module takes care of serving the newly trained model to check for any data drift. If a new distribution of input data is identified (e.g., content generated with a new technique), it is immediately saved and flagged so that the model can be verified and possibly retrained on the new data.

4. Experiments

In this section, we discuss the analyses conducted to evaluate the effectiveness of the continual learning methods introduced

in Section 3.1. We begin by introducing the dataset and selected architectures in Section 4.1, and then, in Section 4.2 we analyze the performance of the two continuous learning strategies by comparing them with transfer learning. [The code base will be made available at https://github.com/francescotss/MLOpsDeepfakeDetection.](https://github.com/francescotss/MLOpsDeepfakeDetection)

4.1. Experimental setting

Dataset

We use the CDDDB (Li et al. (2023)) dataset for all our experiments. The dataset offers three different evaluation scenarios: an easy task sequence, a hard task sequence, and a long one. The dataset contains media generated with 5 different types of GAN-based generative models (StyleGAN, BigGAN, CycleGAN, GauGAN, and StarGAN), 5 non-GAN models (Glow, CRN, IMLE, SAN, and FaceForensics++), and two datasets whose origin is unknown (WhichFaceIsReal and WildDeepfake).

For this study, we report the results on the easy and long sequences. We refer to these two sequences as *Easy* and *Long*, respectively.

- The *Easy set* (composed of GauGAN, BigGAN, CycleGAN, IMLE, FaceForensic++, CRN, and WildDeepfake) is used to study the basic behavior of evaluated methods when they address similar generative techniques.
- The *Long set* (composed of GauGAN, BigGAN, CycleGAN, IMLE, FaceForensic++, CRN, WildDeepfake, Glow, StarGAN, StyleGAN, WhichFaceIsReal, and SAN) is designed to encourage methods to better handle long sequences of deepfake detection tasks, where the catastrophic forgetting might become more serious.

For both sequences, we use the official training, validation, and test splits proposed by Li et al. (2023). Additionally, we also extend the original easy and long sets with two extra classes generated from the Stable Diffusion 1.5 dataset generated from Papa et al. (2023) and DeepFloyd-IF images extracted from Elsa D3⁵ to the *Long set*. Each new class consists of 3000, 1000, and 2000 images for training, validation, and testing, respectively. We call these two modified sets the *Easy+* set and the *Long+* set.

The original dataset also reports a *Hard set*, i.e. a set in which the same classes are rearranged to make the sequence particularly complex for a continuous learning model. We do not report the tables with the detailed experiments on this sequence because, with the modifications made on the *Easy+* and *Long+* sets, we are already analyzing the model’s performance when classes very different from those seen previously in the sequence come into play. In any case, we will briefly comment the performances of this set on Section 4.3.

⁵Source: https://huggingface.co/datasets/elisaEU/ELSA_D3/viewer




Fig. 2: Zero-shot performance of the models when trained on GauGAN and tested on the *Long+* set. All models fail to generalize and register a significant drawdown when tested on new fake samples created from new generative methods. Furthermore, we can see that when exposed to families of models similar to the one seen in training, the models perform slightly better. This could be because the image generation process influences the detector’s performance.

Architecture

For our analysis, we selected four different state-of-the-art that have demonstrated to achieve good results on this task (Gragnaniello et al. (2021); Papa et al. (2023)): Resnet-50, Resnet-18, Mobilenet-V2, and Xception.

4.2. Analysis

We now turn to analyze the performance of continuous learning methods with selected backbones. We start by illustrating the results on the *short* set and then extend the considerations to the longer (i.e., the *long* set). We trained each model using early stopping with a patience of 35 up to 250 epochs. We used the Stochastic Gradient Descent (SGD) optimizer to modify the models’ weights, with an initial learning rate of 0.005 and a momentum of 0.1. The learning rate is controlled by a cosine annealing scheduler with a minimum value of 10^{-5} . For pre-processing, we applied a random cropped on each image with a resolution of 128×128 .

Zero shot

We begin by analyzing the results of all models trained on the GauGAN task and evaluating the whole *long* set for a complete benchmark over a heterogeneous set of generative models. As shown in Figure 2, all the models almost always fail to detect tasks outside their training data, indicating a clear lack of generalizability. These evaluations confirm that a model trained on a particular generative technique struggles to detect other types of fake images. In fact, we can see that the model manages to achieve more or less satisfactory performances on media generated with GANs (in particular BigGAN, CycleGAN, and StarGAN), which evidently have characteristics more similar to those seen in the training phase, but it ultimately fails the tasks more complex ones like FaceForensics++, WildDeepfake, Stable Diffusion 1.5 or DeepFloyd-IF.

Easy and Easy+ set

In Table 1, we report the experiments conducted on this set. The table reports the performance of each dataset trained on the entire sequence. For example, Resnet-50 obtains an accuracy of 57.80% on GauGAN after being trained with KD up to the last available dataset in the sequence (i.e., WildDeepfake). From the table, we can draw some initial insights. Starting with the backbones, the Resnet-50 and Mobilenet-V2 achieve the best results on average. As for learning techniques, we can see a general improvement in the performance of continual learning methods compared to transfer learning. In particular, Knowledge Distillation achieves the best performance when combined with Resnet-50. In this specific configuration, the model achieves excellent performance on IMLE (97.30%) and CRN (94.36%), and good performance on CycleGAN (79.85%), while the results on the other datasets range between 51.22% and 64.45%. The results are not surprisingly high, but to understand why, it is necessary to reason about the different types of datasets contained in this sequence. IMLE and CRN both contain media drawn from computer games, which are indeed more easily recognized. GauGAN, BigGAN, and CycleGAN are all datasets generated with Generative Adversarial Networks, so they have similar characteristics. Finally, FaceForensics++ and WildDeepfake are both challenging sets containing diverse generative techniques. These characteristics make the sequence highly varied and, therefore, complex to be classified uniformly well. In particular, as we iterated through the various tasks in the sequence, we noticed that performance can fluctuate significantly depending on the order in which these datasets are used. Compared to state-of-the-art methods, the KD model is outperformed by iCaRL (Rebuffi et al. (2017)) but works better than gradient-based methods like NSCIL (Wang et al. (2021)). However, it is important to highlight that iCaRL requires examples from previous tasks during training, while knowledge distillation only sees examples belonging to the current task.

Method	Model	Type							Average
		GauGAN	BigGAN	CycleGAN	IMLE	FaceForensics++	CRN	WildDeepfake	
Transfer learning	ResNet-50	57.80	54.38	58.79	<u>56.89</u>	<u>71.63</u>	59.32	71.02	61.40
	ResNet-18	48.70	51.50	44.87	65.02	58.23	<u>66.25</u>	67.32	57.41
	Mobilenet-V2	<u>54.30</u>	<u>49.38</u>	<u>54.40</u>	72.03	74.88	78.82	73.31	65.30
	Xception	45.90	51.38	45.79	48.87	59.81	46.16	65.26	51.88
Class-incremental Learning	NSCIL (Wang et al. (2021))	48.35	50.25	49.43	56.58	56.56	70.73	57.63	55.65
	iCaRL Resnet-32 (Rebuffi et al. (2017))	76.90	80.00	88.93	99.41	85.03	99.45	76.64	87.05
EWC	ResNet-50	65.30	61.13	62.64	68.11	<u>73.77</u>	<u>71.77</u>	70.24	67.56
	ResNet-18	54.15	48.50	58.61	<u>72.50</u>	73.21	69.07	71.36	63.91
	Mobilenet-V2	<u>54.75</u>	48.88	63.37	76.96	86.98	89.62	<u>68.03</u>	69.80
	Xception	50.00	48.75	50.00	56.85	59.72	58.30	61.06	54.95
KD	ResNet-50	64.45	58.75	79.85	<u>97.30</u>	<u>64.65</u>	<u>94.36</u>	51.22	72.94
	ResNet-18	50.50	49.13	52.75	52.27	52.28	53.45	49.87	51.46
	Mobilenet-V2	<u>54.40</u>	<u>53.00</u>	<u>54.21</u>	80.28	72.19	86.37	<u>71.36</u>	67.40
	Xception	50.75	50.75	50.55	97.73	49.86	96.28	49.57	63.64

Table 1: Transfer learning, EWC, and KD performance on the *Easy* set. The average column represents the average accuracy of each model across all datasets. Bold and underlined values represent the highest value and second best for each learning method.

Specifically, given a new task D_t at time t , iCaRL requires the sequence $\mathcal{D} = \{D_1, \dots, D_t\}$ to be trained continuously.

In Figures 3 and 4, we elaborate more on this behavior. The figure shows the average accuracy on each dataset recorded by training the model on the various tasks in the following order: (1) GauGAN, (2) BigGAN, (3) CycleGAN, (4) IMLE, (5) FaceForensics++, (6) CRN, (7) WildDeepfake, (8) Stable Diffusion 1.5, and (9) DeepFloyd-IF. From the figure, we can see that the performance remains high on average on the first three (GAN-based) tasks and then undergoes an initial slight decrease with IMLE and stabilizes with a more substantial decrease from the fifth task onward. By analyzing this behavior, we can infer that the big difference between the generative techniques present in FaceForensic++ and the previous tasks strains the model in finding a region of optimum that reduces the error on all tasks. This result is further confirmed with the arrival of WildDeepfake, Stable Diffusion 1.5, and DeepFloyd-IF, which turn out to be the most complex and different tasks in the sequence. The strong *asymmetry* of some tasks compared to others seems to play an important role, leading the model to optimize performance towards some families of tasks rather than others.

To confirm this hypothesis, we repeated the experiment by removing WildDeepfake from the easy set. As we can see from Figure 5, the performance of the models improves significantly if we compare the performance of the model trained on the complete sequence (Figure 5a) compared to the one trained on the sequence without WildDeepfake (Figure 5b). This tells us that the symmetry between the different tasks is fundamental to maintaining overall satisfactory performance on the entire sequence.

In Table 2, we measure the robustness of all models trained with KD to three image transformations as done in Cocchi et al. (2023): (1) brightness, (2) contrast, and (3) JPEG compression. As shown in the table, all models register a minor degradation of their performance, although Resnet-50 confirms being the best method followed by the Mobilenet-V2.

Overall, the results suggest a few things. First, we can notice a greater stability of continuous learning techniques compared to transfer learning. Furthermore, Knowledge Distillation appears to be more robust overall than EWC, achieving higher performance on average than its rival. In all cases, all techniques achieve better results than the zero-shot scenario. How-

ever, the order of the sequence and the similarity of the tasks seem to play a significant role in overall performance. Finally, the most robust model seems to be the Resnet-50, followed closely by the Mobilenet-V2.

Long set

To complete the analysis, we extend the considerations made up to now to the *long* sequence. In Figure 7a, we report the average accuracy values on each task. As for the *short* sequence, the heterogeneity of the 12 tasks poses a challenge to the continual learning methods, preventing them from learning a general representation applicable to all tasks. In particular, the similar result between the continual learning methods and transfer learning suggests that the latter cannot mitigate *catastrophic forgetting*. In fact, we can note that the models are able to achieve good performance on some tasks like IMLE and CRN (which have similar characteristics) and GAN-based tasks like GauGAN, BigGAN, StarGAN, and StyleGAN but fail to maintain acceptable performance on the other tasks.

Multi-task sequences

Collecting the considerations made so far, in this last test, we test the methods on a hybrid continuous learning scenario. Instead of learning each individual task separately, we combined the *Long* sequence tasks into groups of three, which we call *Multi-tasks*. Therefore, in this configuration, the sequence is composed of 4 macro tasks:

- $t_1 = \{GauGAN, BigGAN, CycleGAN\}$;
- $t_2 = \{IMLE, FaceForensics, CRN\}$;
- $t_3 = \{WildDeepfake, Glow, StarGAN\}$;
- $t_4 = \{StyleGAN, WhichFaceIsReal, SAN\}$.

There is no particular reason behind these four groups. We selected them following the original order of the long set. In this experiment, we are, in fact, interested in studying the behavior of the model in the presence of different task batches.

Figure 7b reports the average accuracy values on each task. As was reasonable to expect, by increasing the heterogeneity of the training data for each task, the model records slightly more robust and uniform performances between the different

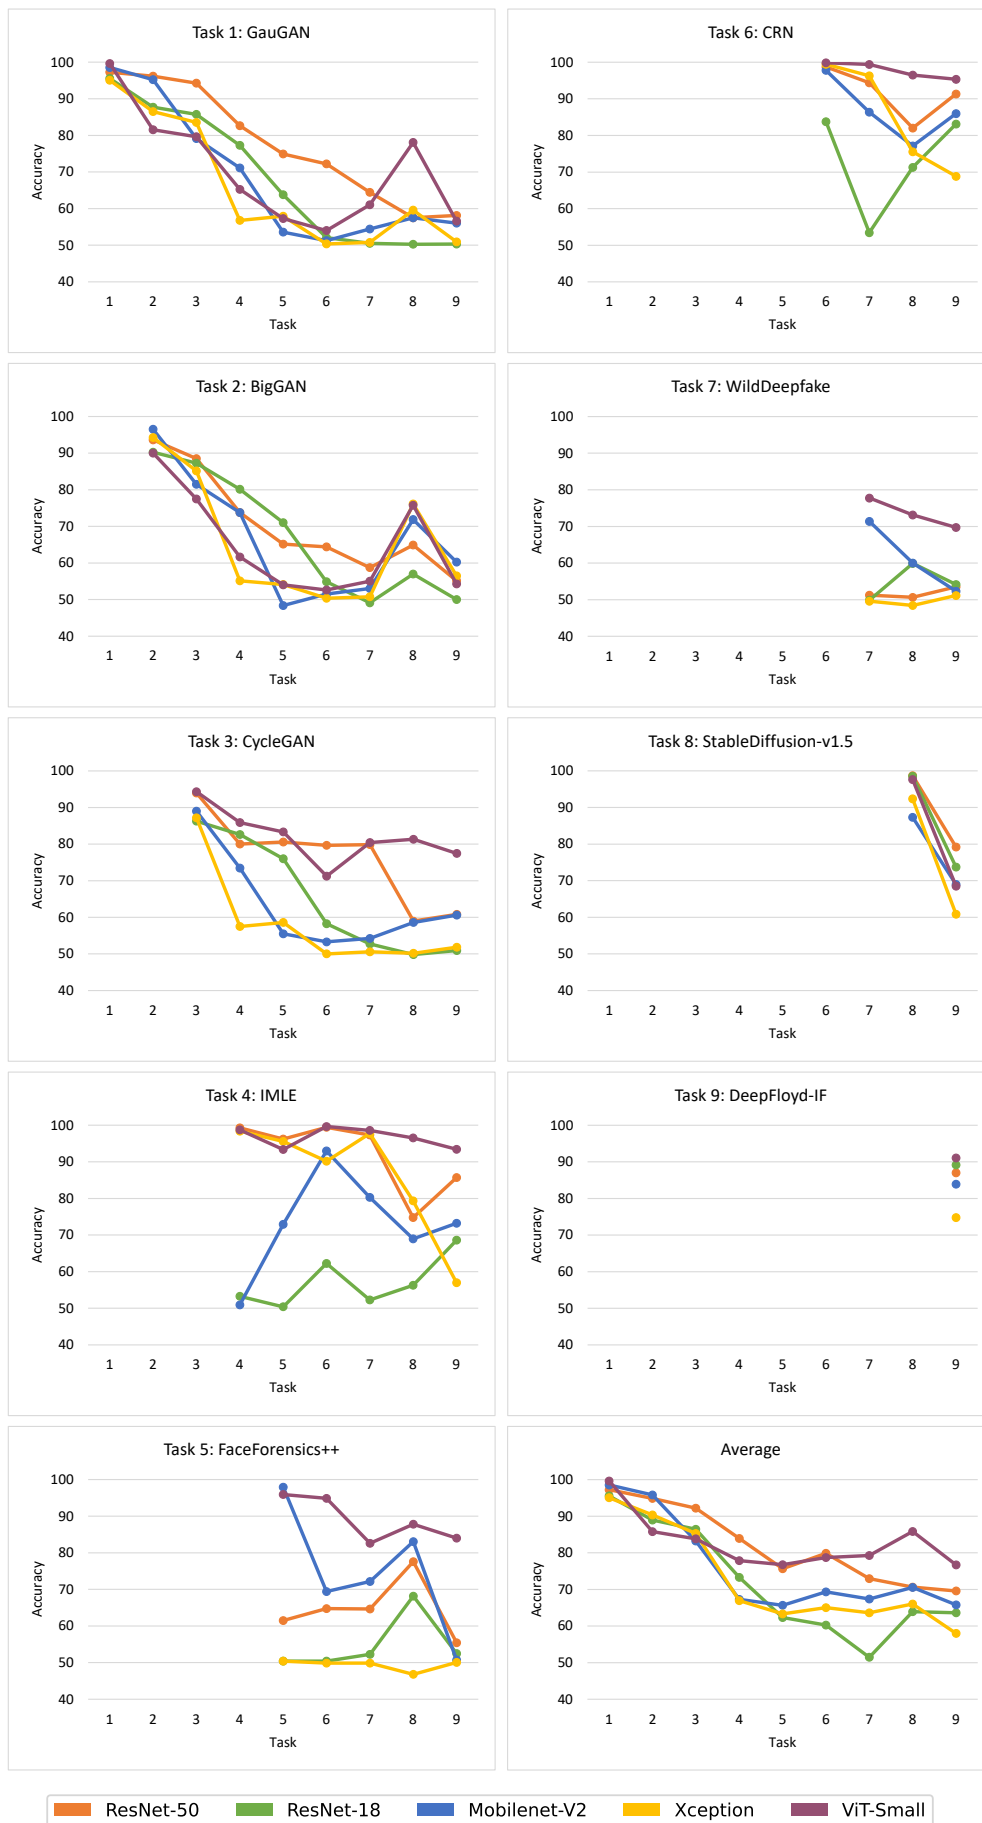


Fig. 3: Knowledge distillation average accuracy at each task t calculated over tasks $\{1, \dots, t\}$ of all backbones on the *Easy+* set. The order of the tasks is the following: (1) GauGAN, (2) BigGAN, (3) CycleGAN, (4) IMLE, (5) FaceForensics++, (6) CRN, (7) WildDeepfake, (8) Stable Diffusion 1.5, and (9) DeepFloyd-IF. The last plot shows the average accuracy across all classes.

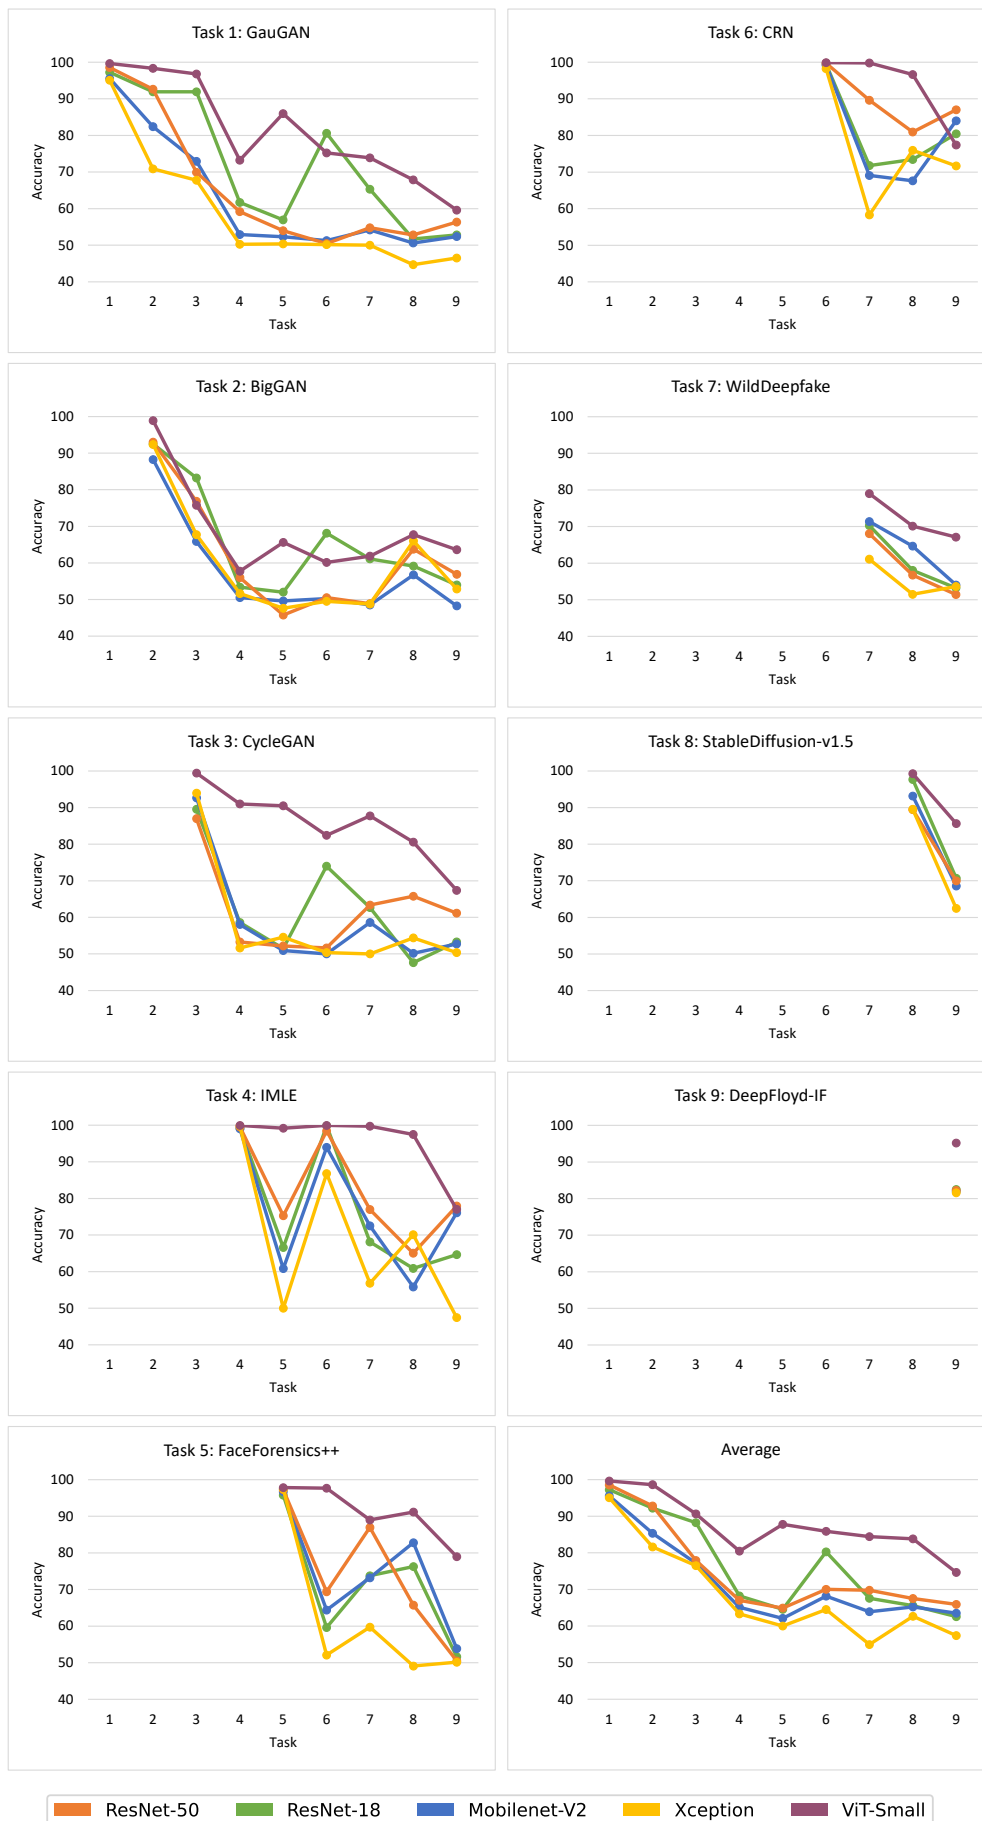




Fig. 4: Elastic Weight Consolidation average accuracy at each task t calculated over tasks $\{1, \dots, t\}$, of all backbones on the *Easy+* set. The order of the tasks is the following: (1) GauGAN, (2) BigGAN, (3) CycleGAN, (4) IMLE, (5) FaceForensics++, (6) CRN, (7) WildDeepfake, (8) Stable Diffusion 1.5, and (9) DeepFloyd-IF. The last plot shows the average accuracy across all classes.

Model	Transformation	Type							Average
		GauGAN	BigGAN	CycleGAN	IMLE	FaceForensics++	CRN	WildDeepfake	
ResNet-50	Brightness	58.85	54.75	70.15	90.45	57.49	86.88	50.14	66.96
	Contrast	59.70	56.75	70.70	94.91	59.35	90.21	51.14	68.84
	JPEG Compression	69.30	60.75	74.18	83.92	58.51	80.93	51.22	68.40
ResNet-18	Brightness	55.30	56.00	60.26	58.57	61.3	58.18	69.45	59.87
	Contrast	54.90	56.50	64.10	57.32	58.98	56.85	71.55	60.03
	JPEG Compression	54.55	54.37	58.97	56.53	57.58	53.09	64.21	57.05
Mobilenet-V2	Brightness	54.2	50.88	55.68	76.56	66.14	82.18	68.36	64.86
	Contrast	54.35	50.88	55.13	76.92	72.74	84.42	67.28	65.96
	JPEG Compression	51.45	51.63	47.25	61.19	56.37	62.80	53.43	54.87
Xception	Brightness	50.35	50.13	49.45	84.27	49.86	85.75	49.61	59.92
	Contrast	50.90	51.25	51.47	83.92	50.14	82.97	49.34	60.00
	JPEG Compression	52.40	48.88	50.55	78.56	50.88	77.45	48.75	<u>58.21</u>

Table 2: Accuracy performance on the *Easy* set with *Brightness*, *Contrast*, and *JPEG Compression* transformations applied using KD models. The average column represents the average performance of each model across all datasets. Bold and underlined values represent the highest value and second best.



(a) Complete *Easy* set.



(b) *Easy* set without WildDeepfake.

Fig. 5: The average accuracy of Resnet-50 trained with KD on the full *Easy* set (Figure 5a) and without WildDeepfake (Figure 5b). Some datasets seem to heavily afflict performance in the continuous learning context.

tasks. Moreover, from the figure we can immediately see that, as suggested in the previous experiment, WildDeepfake puts a strain on all learning techniques. It becomes clear that including images from unknown sources and tasks with a restricted dataset causes a drastic drop in accuracy. Moreover, regarding the learning techniques, all methods achieved satisfactory results. In particular, Knowledge Distillation confirms itself as more robust in most of the sequence but seems to suffer a slight decline in the last three tasks. This result confirms that by aggregating tasks on the basis of their similarity, continuous learning techniques are able to maintain good performance over time.

Model	Metric	Type						
		GauGAN	BigGAN	CycleGAN	IMLE	FaceForensics++	CRN	WildDeepfake
ResNet-50	Precision	96.90	90.63	97.63	99.42	70.14	99.65	52.29
	Recall	28.10	21.75	60.44	93.66	48.03	89.51	21.32
	F1-score	43.57	35.08	74.66	96.45	57.02	94.31	30.29
ResNet-18	Precision	58.53	57.28	63.36	72.18	61.21	70.46	72.56
	Recall	31.90	44.25	67.77	28.01	75.80	28.58	75.54
	F1-score	41.29	49.93	65.49	40.36	67.73	40.67	74.02
Mobilenet-V2	Precision	62.89	55.49	57.14	81.72	75.55	83.23	70.96
	Recall	22.20	25.25	57.14	76.60	70.73	90.52	73.82
	F1-score	32.82	34.71	57.14	79.08	73.06	86.72	72.36
Xception	Precision	81.25	35.71	81.82	97.58	00.00	97.01	43.69
	Recall	01.30	01.25	03.30	97.97	00.00	94.13	03.37
	F1-score	02.56	02.42	06.34	97.77	00.00	95.55	06.25

Table 3: Precision, recall, and F1-score on the *Easy* set using KD models.

4.3. Limitations

The analyzed results give us some interesting information. First of all, continuous learning techniques show an improvement compared to the zero-shot scenario. However, the similarity of the various tasks in sequence and their order of arrival seem to play an important role in maintaining satisfactory performance throughout the sequence. When trained on the *Hard* set, the model struggles to maintain stable performances (GauGAN 60.00%, BigGAN 54.12%, WildDeepfake 49.49%, WhichFaceReal 54.25%, San 82.95%), registering an average 60.16% accuracy.

The results presented require further investigation. First of all, in this study, we limited ourselves to analyzing the behavior of 5 backbones commonly used in deepfake detection; however, a comparison of more advanced deepfake detection methods present in the literature is necessary. Furthermore, the results obtained in the multi-task configuration suggest that these






Fig. 6: Validation accuracies for each task in the *Easy+* set during the training on Task 9 (DeepFloyd-IF). The training task (9) is highlighted in bold.



(a) Average accuracy for Resnet-50.



(b) Average accuracy for Resnet-50 trained with the multi-task configuration.

Fig. 7: Average accuracy on the *Long* set.

techniques could benefit from using a memory (del Rio et al. (2023)). This aspect will be further analyzed in our future studies. Finally, in this study, we assumed the different tasks arrived in batches. This is certainly possible, as the release of a new tool could lead to introducing a significant sample of me-

dia generated with it. However, it remains essential to study the robustness of these techniques on smaller and more heterogeneous sequences.

5. Conclusion

In this paper, we proposed an analysis of two deepfake detection techniques (Knowledge Distillation and Elastic Weight Consolidation), comparing them with zero-shot scenarios and transfer learning. The results show that continuous learning techniques can help make deepfake detection models more robust and easily updated to new generative methods. However, our analysis also highlighted some problems. The performance of these learning strategies seems to depend significantly on the similarity of the tasks and their order of arrival. To address this problem, we have shown how it is possible to combine the different tasks to obtain significantly better performance. Consequently, future developments of this work could analyze the importance of using memory and the robustness of these learning techniques to smaller and more heterogeneous sequences.

In addition to this, we also gave an overview of a CI/CD pipeline for deepfake detection, showing how the models used in this work can be combined with other modules to obtain a pipeline that can be used in a real application scenario. In this regard, in the future we will present a complete version of all the fundamental modules of this pipeline, starting first from the data drift detection module. This module, in particular, could help to significantly improve performance in continuous learning.

Acknowledgments

This study has been partially supported by SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU and Sapienza University of Rome project “EV2” (003_009_22).

References

Aghasani, A., Kangin, D., Angelov, P., 2023. Interpretable-through-prototypes deepfake detection for diffusion models, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 467–474.

- Amerini, I., Anagnostopoulos, A., Maiano, L., Celsi, L.R., 2021. Deep learning for multimedia forensics. *Foundations and Trends in Computer Graphics and Vision* 12, 309 – 457. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85113781771&doi=10.1561/2f0600000096&partnerID=40&md5=65351c1ce97c60be00e45461153b18e7>, doi:10.1561/0600000096.
- Amoroso, R., Morelli, D., Cornia, M., Baraldi, L., Del Bimbo, A., Cucchiara, R., 2023. Parents and children: Distinguishing multimodal deepfakes from natural images. *arXiv preprint arXiv:2304.00500*.
- Chai, L., Bau, D., Lim, S.N., Isola, P., 2020. What makes fake images detectable? understanding properties that generalize, in: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, Springer. pp. 103–120.
- Ciamarra, A., Caldelli, R., Becattini, F., Seidenari, L., Del Bimbo, A., 2024. Deepfake detection by exploiting surface anomalies: the surfake approach, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1024–1033.
- Cocchi, F., Baraldi, L., Poppi, S., Cornia, M., Baraldi, L., Cucchiara, R., 2023. Unveiling the impact of image transformations on deepfake detection: An experimental analysis, in: *International Conference on Image Analysis and Processing*, Springer. pp. 345–356.
- Corvi, R., Cozzolino, D., Poggi, G., Nagano, K., Verdoliva, L., 2023. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 973–982.
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T., 2022. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 3366 – 3385. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-851100848406&doi=10.1109/2fTPAMI.2021.3057446&partnerID=40&md5=678a8e0adef830ba7dd7e5d2ce74fd3>, doi:10.1109/TPAMI.2021.3057446.
- Epstein, D.C., Jain, I., Wang, O., Zhang, R., 2023. Online detection of ai-generated images, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 382–392.
- Fei, J., Xia, Z., Tondi, B., Barni, M., 2023. Robust retraining-free gan fingerprinting via personalized normalization, in: *2023 IEEE International Workshop on Information Forensics and Security (WIFS)*, IEEE. pp. 1–6.
- Fernandez, P., Sablayrolles, A., Furon, T., Jégou, H., Douze, M., 2022. Watermarking images in self-supervised latent spaces, in: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE. pp. 3054–3058.
- French, R.M., 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* 3, 128 – 135. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0032923221&doi=10.1016/2fS1364-6613%2899%2901294-2&partnerID=40&md5=eefdbb22439339cb61278baf7bfaa715>, doi:10.1016/S1364-6613(99)01294-2.
- Gragnaniello, D., Cozzolino, D., Marra, F., Poggi, G., Verdoliva, L., 2021. Are gan generated images easy to detect? a critical analysis of the state-of-the-art, in: *2021 IEEE international conference on multimedia and expo (ICME)*, IEEE. pp. 1–6.
- Hewage, N., Meedeniya, D., . Machine learning operations: A survey on mlops tool support. *arXiv 2022*. *arXiv preprint arXiv:2202.10169*.
- Hinton, G., Vinyals, O., Dean, J., 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Ju, Y., Jia, S., Ke, L., Xue, H., Nagano, K., Lyu, S., 2022. Fusing global and local features for generalized ai-synthesized image detection, in: *2022 IEEE International Conference on Image Processing (ICIP)*, IEEE. pp. 3465–3469.
- Khan, S.A., Dai, H., 2021. Video transformer for deepfake detection with incremental learning, in: *Proceedings of the 29th ACM International Conference on Multimedia*, pp. 1821–1828.
- Kim, M., Tariq, S., Woo, S.S., 2021. Cored: Generalizing fake media detection with continual representation using distillation, in: *Proceedings of the 29th ACM International Conference on Multimedia*, pp. 337–346.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R., 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America* 114, 3521 – 3526. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85016395012&doi=10.1073/2fpnas.1611835114&partnerID=40&md5=e3f68d57af5b09794574205c1b7c321a>, doi:10.1073/pnas.1611835114.
- Kreuzberger, D., Kühl, N., Hirschl, S., 2023. Machine learning operations (mlops): Overview, definition, and architecture. *IEEE access* 11, 31866–31879.
- Le, B.M., Woo, S.S., 2023. Quality-agnostic deepfake detection with intra-model collaborative learning, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 22378–22389.
- Leite, L., Rocha, C., Kon, F., Milojicic, D., Meirelles, P., 2019. A survey of devops concepts and challenges. *ACM Computing Surveys (CSUR)* 52, 1–35.
- Li, C., Huang, Z., Paudel, D.P., Wang, Y., Shahbazi, M., Hong, X., Van Gool, L., 2023. A continual deepfake detection benchmark: Dataset, methods, and essentials, p. 1339 – 1349. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85149024582&doi=10.1109/2fWACV56688.2023.00139&partnerID=40&md5=5a5e6cabff079587edcc50a8ab8da871>, doi:10.1109/WACV56688.2023.00139.
- Lu, Y., Ebrahimi, T., 2024. Assessment framework for deepfake detection in real-world situations. *EURASIP Journal on Image and Video Processing* 2024, 6.
- Maiano, L., Papa, L., Vocaj, K., Amerini, I., 2022. Depthfake: a depth-based strategy for detecting deepfake videos. *arXiv preprint arXiv:2208.11074*.
- Marra, F., Saltori, C., Boato, G., Verdoliva, L., 2019. Incremental learning for the detection and classification of gan-generated images, in: *2019 IEEE international workshop on information forensics and security (WIFS)*, IEEE. pp. 1–6.
- Ojha, U., Li, Y., Lee, Y.J., 2023. Towards universal fake image detectors that generalize across generative models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24480–24489.
- Paleyas, A., Urma, R.G., Lawrence, N.D., 2022. Challenges in deploying machine learning: A survey of case studies. *ACM Computing Surveys* 55. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85146490462&doi=10.1145/2f3533378&partnerID=40&md5=9927d80efde171e84d327c1e994d4017>, doi:10.1145/3533378.
- Pan, K., Yin, Y., Wei, Y., Lin, F., Ba, Z., Liu, Z., Wang, Z., Cavallaro, L., Ren, K., 2023. Dfil: Deepfake incremental learning by exploiting domain-invariant forgery clues, p. 8035 – 8046. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85179547005&doi=10.1145/2f3581783.3612377&partnerID=40&md5=73efaa1437d0f5f9b89df8e3ebf5c66f>, doi:10.1145/3581783.3612377.
- Papa, L., Faiella, L., Corvito, L., Maiano, L., Amerini, I., 2023. On the use of stable diffusion for creating realistic faces: from generation to detection, in: *2023 11th International Workshop on Biometrics and Forensics (IWBF)*, IEEE. pp. 1–6.
- Raza, M.A., Malik, K.M., 2023. Multimodaltrace: Deepfake detection using audiovisual representation learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 993–1000.
- Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H., 2017. icarl: Incremental classifier and representation learning, in: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010.
- del Rio, F., Hurtado, J., Buc, C., Soto, A., Lomonaco, V., 2023. Studying generalization on memory-based methods in continual learning. *arXiv preprint arXiv:2306.09890*.
- Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., Niessner, M., 2019. Faceforensics++: Learning to detect manipulated facial images, p. 1 – 11. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85081924374&doi=10.1109/2fICCV.2019.00009&partnerID=40&md5=0d3b9a49170ab85ca82d1b415fdb5532>, doi:10.1109/ICCV.2019.00009.
- Semola, R., Lomonaco, V., Bacchi, D., 2022. Continual-learning-as-a-service (claas): On-demand efficient adaptation of predictive models. *arXiv preprint arXiv:2206.06957*.
- Sha, Z., Li, Z., Yu, N., Zhang, Y., 2023. De-fake: Detection and attribution of fake images generated by text-to-image generation models, in: *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pp. 3418–3432.
- Verdoliva, L., 2020. Media forensics and deepfakes: An overview. *IEEE Journal on Selected Topics in Signal Processing* 14, 910 –

932. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85088105202&doi=10.1109%2fJSTSP.2020.3002101&partnerID=40&md5=295a13ed94c814b67cda750ad3a1796e>, doi:10.1109/JSTSP.2020.3002101.
- Wang, J., Alamayreh, O., Tondi, B., Barni, M., 2023. Open set classification of gan-based image manipulations via a vit-based hybrid architecture, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 953–962.
- Wang, S., Li, X., Sun, J., Xu, Z., 2021. Training networks in null space of feature covariance for continual learning. URL: <https://arxiv.org/abs/2103.07113>, arXiv:2103.07113.
- Wang, S.Y., Wang, O., Zhang, R., Owens, A., Efros, A.A., 2020. Cnn-generated images are surprisingly easy to spot... for now, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 8695--8704.
- Wu, H., Liu, G., Yao, Y., Zhang, X., 2020. Watermarking neural networks with watermarked images. IEEE Transactions on Circuits and Systems for Video Technology 31, 2591--2601.
- Zhang, Y., Lin, W., Xu, J., 2024. Joint audio-visual attention with contrastive learning for more general deepfake detection. ACM Transactions on Multimedia Computing, Communications and Applications 20, 1--23.
- Zhao, Y., Liu, B., Ding, M., Liu, B., Zhu, T., Yu, X., 2023. Proactive deepfake defence via identity watermarking, in: Proceedings of the IEEE/CVF winter conference on applications of computer vision, pp. 4602--4611.
- Zi, B., Chang, M., Chen, J., Ma, X., Jiang, Y.G., 2020. Wilddeepfake: A challenging real-world dataset for deepfake detection, p. 2382 { 2390. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85099811365&doi=10.1145%2f3394171.3413769&partnerID=40&md5=a89cb90a359fde3afefeff6cd66d3ced>, doi:10.1145/3394171.3413769.