Learning to Learn Face-PAD: a lifelong learning approach

Daniel Pérez-Cabo*
Gradiant - UVigo, Spain
dpcabo@gradiant.org

David Jiménez-Cabello
Nielsen Connect R&D, Spain
david.jimenez.phd@gmail.com

Artur Costa-Pazo
ALiCE Biometrics - UVigo, Spain
acosta@alicebiometrics.com

Roberto J. López-Sastre
University of Alcalá, Spain
robertoj.lopez@uah.es

Abstract

A face presentation attack detection (face-PAD) system is in charge of determining whether a face corresponds to a presentation attack or not. The vast majority of proposed solutions consider a static scenario, where models are trained and evaluated in datasets where all types of attacks and conditions are known beforehand. However, in a real-world scenario, the situation is very different. There, for instance, the types of attacks change over time, with new impersonation situations appearing for which little training data is available. In this paper we propose to tackle these problems presenting for the first time a continual learning framework for PAD. We introduce a continual meta-learning PAD solution that can be trained on new attack scenarios, following the continual few-shot learning paradigm, where the model uses only a small number of training samples. We also provide a thorough experimental evaluation using the GRAD-GPAD benchmark. Our results confirm the benefits of applying a continual meta-learning model to the real-world PAD scenario. Interestingly, the accuracy of our solution, which is continuously trained, where data from new attacks arrive sequentially, is capable of recovering the accuracy achieved by a traditional solution that has all the data from all possible attacks from the beginning. In addition, our experiments show that when these traditional PAD solutions are trained on new attacks, using a standard fine-tuning process, they suffer from catastrophic forgetting while our model does not.

1. Introduction

Face presentation attack detection (face-PAD) is indeed a crucial and challenging problem for face recognition commercial solutions. Technically, a face-PAD solution must determine whether the face presented to the system is a bona fide presentation (BFP) or a presentation attack (PA). This functionality has a tremendous impact in any commercial service based on face recognition technology (*e.g.* access control in airports or borders, phone unlock, *etc.*).

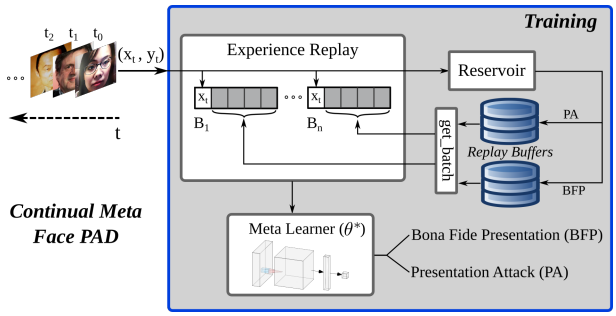

Figure 1: Graphical abstract of the proposed continual meta face anti-spoofing approach (CM-PAD).

A wide variety of mediums are used to perform presentation attacks (*e.g.* videos, printed pictures, masks, *etc.*). But, attackers will always be looking for new ways to bypass the security of facial recognition systems. This situation forces us to explore new solutions that can learn to detect new types of attacks, with a clear restriction: the lack of training data for them.

Current state-of-the-art PAD approaches have achieved unprecedented results in standard benchmarks and datasets [22, 26]. The reasons are clear: a) available data is orders of magnitude under model's parameters; and b) all the attacks are known beforehand and are completely represented by the training data. The problem arises when the trained PAD models have to operate in a real-world scenario, dealing with unseen attacks, for instance. Here the generalization problem arises: state-of-the-art approaches exhibit a severe drop of performance due to the possible overfitting that maximizes their accuracy for just

*Authors are sorted depending on their contribution to the paper.

978-1-7281-9186-7/20/$31.00 @2020 European Union

the benchmarks used to train and validate them [9].

Practical solutions for this problem consist in retraining the AI models completely from scratch or fine tuning the previous models with the novel data. But these approaches present two important weaknesses: 1) the catastrophic forgetting situation is not avoided, where with the retraining process the PAD models lose performance in the previously learned attacks; and 2) the scarcity of training data for new attacks. To directly address these two problems in this paper we propose to model the PAD problem with a continual meta-learning approach.

Our continual meta-learning PAD approach can be trained on new scenarios (*e.g*. attacks, domains, *etc*.) as data appears, without the need to retrain the model from scratch. We follow the continual learning paradigm [23] where the model must learn from an infinite stream of data, with the target of incrementally extending the acquired knowledge but without catastrophic forgetting. This learning strategy imitates the way of working in a realistic scenario such as the one described above. But we also need a meta-learning solution [14], in the sense that our PAD model should be able to solve new scenarios using a small number of training samples.

As we show in Fig. 1 we propose a PAD model that adopts a meta-learning paradigm on the continual learning problem. The model learns to learn each new situation generalizing to other unseen examples, such as new attacks or domains. Therefore, the main contributions of our work are:

- For the first time, a continual learning approach is applied to PAD using meta-learning for fast adaptation. Technically, we build on the recently proposed Meta Experience Replay (MER [29]) model to adapt it for the PAD scenario. We propose to extend the MER model with a double replay-buffer strategy that allocates samples from real face and impostor face categories in a separate way when learning every new attack. This allows us to balance the influence of the two categories during the meta-training steps. In Section 3 we provide all these details.

- In order to provide a thorough experimental validation for the new PAD scenario described, we have carefully designed two novel evaluation protocols for the GRAD-GPAD benchmark [8] (see Section 4.3). The proposed continual learning protocols and baselines will all be made publicly available to encourage (much needed) further research on continual meta-learning PAD.

- The results reveal the benefits of a continual meta-learning PAD model. First, our approach is able to achieve a similar accuracy as the proposed PAD backbone trained offline, despite the sequential availability of training data. Note that ours is also a meta-learning based model, hence able to learn from few samples. Second, our experiments also show that when the traditional PAD solutions are trained on new attacks, following a standard fine-tuning process, they suffer from catastrophic forgetting, while our model exhibits great immunity to this phenomenon. Section 4 includes all the experimental evaluation.

GRAD-GPAD stands for Generalization Representation over Aggregated Datasets for Generalized PAD

https://github.com/dpcabo/CM-PAD-lists

2. Related Work

Face Recognition (FR) systems have become a mature technology to use in real environments, however they can be easily fooled by simple impersonation attacks (*e.g*. Print or Replay). Thus, many face-PAD approaches have been proposed in the literature to provide security to FR systems. We refer the reader to the following two excellent surveys [15, 28], where thorough taxonomies of PAD methods for FR systems are offered.

Whether these PAD techniques can be applied in real world scenarios, has constituted an active line of work in recent years. Generalized presentation attack detection (GPAD) is a problem that has been studied from different perspectives. One of the first works to raise this issue is [10] where the authors reformulate the traditional binary classification problem as an anomaly detection approach using texture-based features. With the rise of deep neural networks, different approaches have been applied: following domain adaptation techniques [20]; learning discriminative features in the embedding space [18, 19, 20]; using generative models [18]; or implementing anomaly detection mechanisms [26]. In [8, 9, 11, 25], we find works that focus on the analysis of the GPAD problem. Overall, these works show that the generalization capability of most state-of-the-art PAD approaches is not enough when they are faced to a realistic scenario. In other words, most of them exhibit a clear overfitting towards the training data [25], when, for instance, they are evaluated following generalization-specific protocols (*e.g*. Cross-Dataset, Unseen Capture Device) in the novel GRAD-GPAD benchmark [8, 9].

However, to the best of our knowledge, there are no works in the literature that address the problem of GPAD combining continual learning and meta-learning. This is actually the approach we introduce in this work.

Meta-learning [2, 14, 24, 27, 30], within the context of image classification, addresses the problems of data scarcity and generalization, trying to mimic the learning process of humans in two aspects: 1) using few samples to learn a new task (*i.e*. few-show learning); and 2) rapidly adapting its knowledge to behave well under new scenarios.

With respect to continual learning, it is defined as the paradigm where the models are able to learn through time

without forgetting. Classical *off-line* supervised deep neural networks struggle on building upon previous experience, and usually rely on training from scratch or fine-tuning the model as new data is collected. The former suffers from the burden of increasing computational effort through time, whereas the latter comes under catastrophic forgetting that results in exponential loss of the retained knowledge. We refer to [12] for an excellent survey on continual learning.

In this paper, we bring for the first time the continual learning paradigm to face-PAD. Our approach is able to mitigate the generalization problems that arise when new attack types or new scenarios appear. We contribute to GRAD-GPAD benchmark [8] by adding an exhaustive and hierarchical categorization that enables its use in both meta-learning and continual learning settings. The experiments carried out show that the proposed offline backbone is able to perform on par to the state-of-the-art. Besides, our continual meta face-PAD approach is able to highly reduce the effect of catastrophic forgetting with almost no penalty, adapting rapidly to unseen scenarios without having to retrain the model from scratch.

3. Continual meta-learning model for face-PAD

This section starts with an introduction to both the continual and the meta-learning paradigms. Once this background has been provided, we proceed to present our proposal to address the face-PAD problem.

3.1. Background on Continual and Meta-Learning

3.1.1 Meta-learning

Optimization-based meta learning considers a function approximator $f_\theta(x)$ that maps input samples x to labels y and a set of tasks $\mathcal{T}_i$ drawn from a distribution over tasks $p(\mathcal{T})$. Note that for the particular context of face-PAD these tasks are the different attacks (PAs) and bona fide presentations (BFPs). One can thus express the corresponding meta-learning optimization objective as:

$$\theta^* = \arg\min_\theta \mathbb{E}_{\mathcal{T}\sim p(\mathcal{T})} \left[\mathcal{L}_\theta(\mathcal{T})\right]. \tag{1}$$

Different from traditional supervised methods, in the meta-learning formulation, we have to consider tasks instead of training pairs (x, y). Each task is generally composed of two disjoint datasets for meta-training and testing, respectively. Different approaches can be found in the literature, to address this meta-learning problem [14, 24]. The authors in [14] propose a model agnostic approach called MAML. It separates meta-learning into two steps. First, MAML inner updates the weights of the model using k steps of gradient descent for m batches of disjoint training tasks. In this step the algorithm learns task-specific weights over a support set $\mathcal{S}$ of training pairs. Once the model updates are obtained ($\phi_\mathcal{T}$), MAML upper updates the final model's parameters by back-propagating the loss due to each task with respect to a set of query samples $\mathcal{Q}$ that wasn't seen so far, i.e. the meta-learning objective implicitly optimizes the generalization loss while training in the same way as testing. The model obtained embeds a parameter initialization that lies closely to the optimal manifold of all tasks and thus it can be optimized for a new task using a reduced set of training pairs. To avoid second order derivatives, the authors propose a first order approach (FOMAML) that obtains similar results.

Following the same intuition as FOMAML, in [24] the authors propose Reptile, a simplification of MAML's objective to avoid second order derivatives. The authors replaced gradients in the outer update of MAML by a simple subtraction between the parameters obtained in the inner loop ($\phi_\mathcal{T}$) and the previous model's weights (θ_{t-1}), optimizing for the following objective:

$$\theta^* = \arg\min_\theta \mathbb{E}_{\mathcal{T}\sim p(\mathcal{T})} \left[\frac{1}{2}\text{dist}\left(\theta, \phi_\mathcal{T}\right)^2\right]. \tag{2}$$

Using a second order Taylor expansion, the authors demonstrate that the expectation over tasks $\mathcal{T}$ depends on the inner product between the gradients of similar tasks providing improved generalization. Given a sequence of n batches ($B_i \sim p(\mathcal{T}),\ \ 0 < i \leq n$) of training pairs, one could rewrite the objective as:

$$\theta^* = \arg\min_\theta \mathbb{E}_{B_1,\cdots,B_n} \left[2\sum_{i=1}^{n}\left[\mathcal{L}(B_i) - \sum_{j=1}^{i-1}\alpha\frac{\partial\mathcal{L}(B_i)}{\partial\theta}\cdot\frac{\partial\mathcal{L}(B_j)}{\partial\theta}\right]\right]. \tag{3}$$

3.1.2 Continual Learning

In the continual learning paradigm an agent is exposed to a non-stationary stream of data shaped like sequential tasks of several observations of the same distribution over tasks $q(\mathcal{T})$. Within the context of face-PAD the continual learning framework introduces the scenario in which systems have to deal with unseen situations (*e.g.* new types of attacks or changing domains), for which little training data is available. The key challenge of continual learning is to avoid catastrophic forgetting while learning incrementally. In other words, and within the context we are dealing with here, we want the face-PAD solutions to be able to adapt to unseen settings, but without losing precision in the previous ones. Thus, in a continual learning model, the overall objective is defined as:

$$\theta_t^* = \arg\min_\theta \mathbb{E}_{x_t\sim q(\mathcal{T})} \left[\mathcal{L}_\theta(x_t)\right], \tag{4}$$

$$\text{s.t.} \quad \mathcal{L}_\theta(x_i) \leq \mathcal{L}_{\theta_{t-1}}(x_i) \quad \forall i \in [0, \ldots, t-1]\,. \tag{5}$$

The constraint imposed by Eq. 5 is known as the elasticity-plasticity dilemma. As suggested by [29], such a constraint can be reformulated in the gradient space as the alignment between the gradients of two arbitrary sample pairs $(x_i, y_i), (x_j, y_j)$:

$$\langle g_i, g_j \rangle = \left\langle \frac{\partial \mathcal{L}(x_i, y_i)}{\partial \theta}, \frac{\partial \mathcal{L}(x_j, y_j)}{\partial \theta} \right\rangle . \tag{6}$$

When the inner product is greater than zero gradients are aligned and transfer occurs, whereas if it is lower than zero, gradients are not aligned and interference appears during training (the transfer-interference trade-off).

3.2. Our approach: Continual Meta Face-PAD

We introduce here a Continual Meta-Learning face-PAD (CM-PAD) approach that leverages the benefits of meta-learning to mitigate the influence of catastrophic forgetting in a continual learning setting. Technically, we build on the recently proposed Meta Experience Replay (MER) model [29]. MER grounds in the intersection of continual and meta-learning. By taking into consideration the elasticity-plasticity dilemma through gradient alignment, MER is able to control the dynamics of weight sharing across time and makes interference less likely to occur while maximizing the chance to learn parameters that transfer future knowledge rapidly. Given two arbitrary training pairs $(x_i, y_i), (x_j, y_j)$, drawn from a distribution over tasks $q(\mathcal{T})$, the overall objective is expressed as:

$$\theta^* = \arg\min_\theta \mathbb{E}_{[x_i, x_j]} \left[\mathcal{L}(x_i, y_i) + \mathcal{L}(x_j, y_j) - \alpha \frac{\partial \mathcal{L}(x_i, y_i)}{\partial \theta} \cdot \frac{\partial \mathcal{L}(x_j, y_j)}{\partial \theta} \right]. \tag{7}$$

In [29] the authors demonstrated that Reptile's formulation from Eq. 3 is a suitable and frictionless framework to integrate meta-learning into MER's formulation when samples are considered sequentially one by one.

To deal with the non-stationary stream of data in continual learning, MER uses experience replay to maintain past knowledge by augmenting learning at every step using a random batch of samples drawn from a replay buffer of past experience with a limited capacity of M samples. They propose to use a Reservoir strategy [31] to replace samples from the replay buffer with probability $\frac{M}{N}$, with N the number of training pairs seen so far.

In the context of a never-ending stream of data where samples are fed one by one to the system, we define face-PAD as a two-class classification problem (PA and BFP) where each task is composed of pairs of training samples belonging to categories that were not considered previously. To ensure that each meta-training step contains samples from both classes, we propose a double replay-buffer that allocates samples from PA and BFP categories in a separate way. Based on the overall distributions of samples within the available datasets, we propose to sample BFPs and PAs from their corresponding replay buffer with the following probabilities: 0.2 for BFPs and 0.8 for PAs. Note that this distribution might change depending on the scenario. In Fig. 2 we show the training procedure of the proposed continual meta-learning framework.

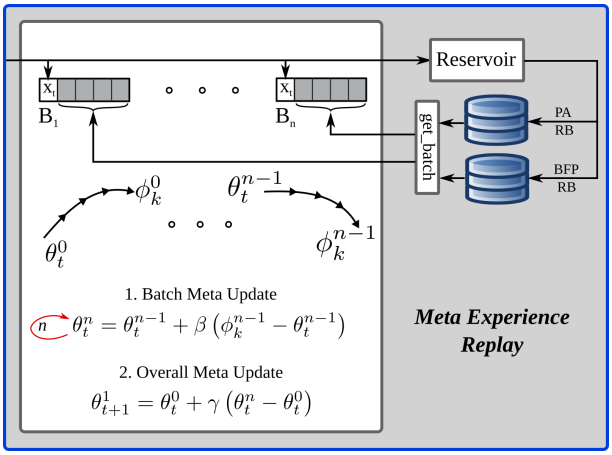

Figure 2: Diagram of the parameter's update for the proposed Continual Meta Face-PAD, where θ_t^i are the batch-specific parameters obtained for task i at instant t, ϕ_k^i refers to θ_t^i after applying k iterations of SGD and β, γ are the gradient step sizes for the batch and overall meta updates, respectively.

3.3. Continual and Meta Learning extension for GRAD-GPAD

In this work we propose to build on [29] to bring for the first time both fields of meta-learning and continual learning together for face-PAD. To this end we extended GRAD-GPAD [8] (the largest aggregated dataset for face-PAD) with new protocols for continual and meta-learning.

Besides, we present an exhaustive and hierarchical categorization that provides a rich set of labels for both BFPs and PAs that suits the meta-learning paradigm. The new protocols for GRAD-GPAD (see Section 4.3) simulate the non-stationary stream of data in form of tasks during training. In Fig. 3 we show the hierarchical tree of categories that allows us to define a rich set of disjoint tasks for learning on a sequential-like training, the Continual Meta Learning Categorization (CMLC).

Based on the common taxonomy proposed in GRAD-GPAD, specific labeling has been created with the following attributes: dataset, device, SPAI (Scene Presentation Attack Instruments, including bona fide accesses) and capture con-

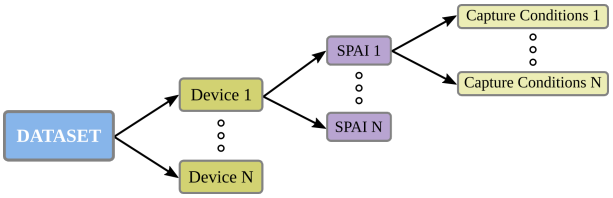

Figure 3: Diagram of the proposed Continual Meta Learning Categorization (CMLC) used to extend GRAD-GPAD for meta-learning settings.

ditions (it combines the different capture conditions represented in each dataset independently).

Table 1 shows the 10 datasets represented in the aggregated GRAD-GPAD framework, where we observe that datasets with greater variability generate a greater number of specific categories and therefore generate a greater number of disjointed tasks. For example, Oulu-NPU [5] has 6 different capture devices and 6 types of PAs, that allows to generate very representative tasks of the fraudulent accesses. On the other hand, Replay-Mobile [7] has only two capture devices, but nevertheless it contains 5 different capture conditions, so it can generate a large number of very useful tasks to represent the conditions of bona fide users in a real case.

4. Experiments

4.1. Experimental setup

Network Architecture and Training Setups Inspired by the recent advances on auxiliary pixel-wise supervision for PAD [1, 22], we propose the architecture in Fig. 4. This model consists of a fully convolutional backbone followed by a depth regression block and one classification block. We add residual connections in the ResNet blocks [17] and, instead of using two different branches for depth regression and classification as in [22], we serialize both of them similar to [1]. The complete architecture of our network is shown in Fig. 4.

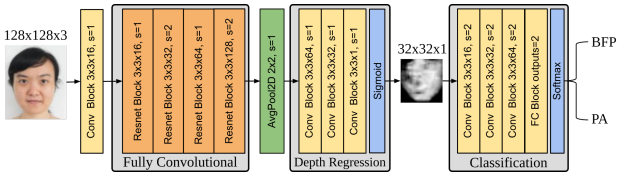

Figure 4: The proposed backbone architecture.

For the experimental evaluation, we propose the next training approaches. First, we train our model following a standard supervised way: we call this model as ***Ours-supervised***. We use 20 epochs in each experiment (unless otherwise stated) and a batch size of 32. See Fig. 4 for the rest of the parameters. The loss function consists of two terms: 1) the *cross-entropy loss* for the final classification output, and 2) the *depth loss* introduced by [22] for the auxiliary depth supervision. Second, we train the deep architecture with our proposal for continual meta-learning (**CM-PAD**). Technically, this approach is trained by feeding samples one by one (note that samples are only considered once during training). We use the following parameters: $\alpha = 0.01$, $\beta = 0.03$, $\gamma = 0.5$, $k = 5$, $n = 20$ and $M = 3000$. For both strategies, we use stochastic gradient descent (SGD) without momentum and Group Normalization [33].

Preprocessing Since we propose a single-frame based approach and most datasets provide videos, we pick the central frame of each video. The input to the network is a cropped face extracted using the MTCNN face detector [34]. We follow the same procedure as in [22] to compute the pseudo ground truth for depth supervision.

Evaluation Metrics To compare our method with prior (non-continual learning) works we use ISO/IEC 30107-3 standard metrics: *i.e.* Attack Presentation Classification Error Rate (APCER), Bonafide Presentation Classification Error Rate (BPCER) and Average Classification Error Rate (ACER). These metrics are designed to evaluate generalization on face PAD problems. We would like to highlight the importance of ACER metric, as it entails the most challenging scenario, where performance is computed for every PA independently, but it only considers the results for the worst case. We argue that the Half Total Error Rate (HTER), one of the most popular metrics for PAD, does not contribute to assess generalization as it is a particular case of ACER where all attacks are considered of the same type.

For the continual meta-learning experiment, we use a specific metric: the Backward Transfer and Interference (BTI), introduced in [29]. This evaluation metric has been designed to measure the effect of the catastrophic forgetting in a continual learning scenario. Technically, BTI is computed as the average difference between the accuracy of each task at the end of the training and when the task was learned. Thus, BTI is able to show the impact of catastrophic forgetting in continual learning settings such that higher negative values mean that the model is more vulnerable to catastrophic forgetting. In all the experiments we use disjoint splits for training and testing.

Previous Works for Comparisons We compare our approach with four state-of-the-art methods. The approach in [25] (*Quality*) computes hand-crafted features based on quality evidences. They obtain a 139-length feature vector from the concatenation of the quality measurements proposed in [16] and [32]. The second method, proposed in [4]

Dataset	Year	# Ids	# Samples real/attack	CMLC real/attack	Spoof attack	# Capture Devices/Displays	# Tasks train/test
CASIA-FASD [35]	2012	30	150/450	3/9	P, R	3/1	9/9
Replay-Attack [6]	2012	50	200/1000	2/6	P, 2xR	1/2	6/6
3DMAD [13]	2013	17	170/85	1/1	M	1/-	1/1
MSU-MFSD [32]	2015	35	110/330	2/6	P, 2xR	2/2	6/6
Replay-Mobile [7]	2016	40	390/640	**10/8**	P, R	2/1	10/10
HKBU [21] (v1)	2016	8	70/40	1/1	M	2/-	1/1
Oulu-NPU [5]	2017	55	1980/3960	**18/72**	2xP, 2xR	6/2	72/72
Rose-Youtu [20]	2018	20	900/2600	2/6	2xP, 2xR, 2xM	5/2	6/6
SiW [18]	2018	165	132/0330	4/8	2xP, 4xR	2/4	8/7
CS-MAD [3]	2018	14	88/220	8/4	P, M	2/-	8/8

Table 1: Overall information of the proposed extension of GRAD-GPAD benchmark for Continual Meta face anti-spoofing. Print, Replay and Mask attacks are represented with P, R and M, respectively.

(*Color*), consists in computing a color-based feature vector of high dimensionality (19998-length) by concatenating texture features based on Local Binary Patterns (LBPs) in two different color spaces (*i.e*. YCbCr and HSV). The third method is the so-called *Auxiliary*, proposed in [22], which introduces a two-branch deep neural network that incorporates pixel-wise auxiliary supervision constrained by the depth reconstruction of the faces (attacks are forced to belong to a plane) and rPPG (*i.e*. remote heart rate monitoring) estimation based on the video frames. In the experiments, we use the same single frame configuration presented in the ablation study of [22]. Finally, in [26] (*Anomaly*) the authors propose to model the problem of PAD from an anomaly detection perspective using metric learning.

Results for [22] and [26] come from our re-implementations. For fair comparisons, all the approaches are trained during 20 epochs, which corresponds with the same number of times that our approach sees the input data (*i.e*. the number of batches for the meta update, $n = 20$).

4.2. Comparison with Previous Works

In this section we compare the proposed backbone and our CM-PAD approach with state-of-the-art models following a traditional supervised training. Regardless this is not a fair comparison for our CM-PAD, we want to assess how far we are from traditional pipelines that, different from us, can iterate through **all the data** several times (20 epochs). For that purpose we use the challenging GRAD-GPAD framework.

In Table 2 we show a comparison with previous works using two traditional supervised protocols (*i.e*. non-continual setting) of GRAD-GPAD: 1) Grandtest and 2) Cross-dataset Test-On-CASIA-FASD, respectively. In *Ours-Supervised*, we provide the performance of our backbone model trained following a 2-class classification problem between PAs and BFPs. Finally, in *Ours-CM-PAD* we show the results using the proposed continual meta-learning setting and the CMLC categorization.

As depicted from Table 2, the proposed backbone (*Ours-Supervised*) performs on par to the best methods. Moreover, we show that introducing the continual meta-learning setting barely penalizes the performance compared to traditional supervised training in the Cross-Dataset experiment. As we might expect, *Ours-CM-PAD* behaves worse than the baseline but yet on par to the other methods. These results are of great relevance in the context of PAD, since we can exploit rapid adaptation and continual learning with almost no impact on global performance and considering only one sample at a time without iterating several epochs through the whole data. This brings the opportunity to quickly adapt to new domains and attacks as data comes in, without sacrificing precision nor needing to retrain the model.

4.3. Continual Meta-Learning Evaluation

In this section we evaluate the benefits of using a continual learning setting, in contrast to a standard stationary learning strategy. For that purpose, we use the continual meta-learning protocols explained in Sec. 4.3.1 and Sec. 4.3.2. We focus the experimental evaluation on the catastrophic forgetting effect.

4.3.1 Evaluation on changing domains

The Continual-Meta-Grandtest protocol. In this protocol we simulate a real face-PAD scenario where we normally move from one domain to another over time, changing, for instance: illumination, capture devices, *etc*. Each of the available datasets in Table 1 is captured using a different setting and thus it represents a new domain with a new set of enriched CMLC categories (see Fig. 3) and their corresponding tasks. For this protocol we sort each dataset sequentially and feed the learning algorithm with their corresponding samples in a continual fashion. We define a Continual-Meta-Grandtest task as a binary classifica-

Algorithm	Grandtest			Cross-Dataset Test-On-CASIA-FASD		
	ACER	BPCER@APCER = 5 %	BPCER@APCER = 10 %	ACER	BPCER@APCER = 5 %	BPCER@APCER = 10 %
Quality [25]	36.99 %	97.75 %	97.49 %	47.38 %	83.15 %	79.77 %
Color [4]	19.21 %	70.31 %	38.79 %	25.69 %	65.73 %	41.57 %
Auxiliary [22]	31.89 %	58.68 %	37.98 %	26.90 %	58.43 %	48.31 %
Anomaly [26]	**10.47 %**	69.85 %	39.75 %	**18.48 %**	**48.31 %**	**28.09 %**
Ours-supervised	14.23 %	**33.76 %**	**18.63 %**	30.56 %	62.22 %	54.44 %
Ours-CM-PAD	28.66 %	49.90 %	37.03 %	32.22 %	64.89 %	54.44 %

Table 2: Comparison with state-of-the-art methods using the Grandtest and the Cross-Dataset Test-On-CASIA-FASD protocols from the GRAD-GPAD framework.

tion problem that differentiates between PAs and BFPs both belonging to the same domain. As a result we obtain a total of 127 binary tasks for this protocol (see CMLC column in Table 1).

Following the new Continual-Meta-Grandtest protocol we show in Fig. 5 a confusion matrix where the values represent the accuracy of the proposed CM-PAD model when tasks are sequentially introduced. In rows we show the tasks sorted in sequential order as they are used for training and, in columns , we show the corresponding evaluation for all the tasks. Thus, the first row in the confusion matrix represents the scenario where we are training *only* for the first task and testing for *all* the tasks in the dataset. Values in the main diagonal correspond to the accuracy of the model when we evaluate in the same task as it is trained on. The points below the main diagonal show the ability of the model to retain knowledge (*i.e*. the catastrophic forgetting). Values above the main diagonal show the capacity of the model to generalize to *unseen scenarios* (*e.g*. new PAs or domains), *i.e*. the zero-shot performance.

In the blue rectangle of Fig. 5, which corresponds to the tasks associated to Oulu-NPU dataset, we observe that our CM-PAD approach is able to perform well for unseen tasks of the same dataset, meaning that it is able to rapidly reuse pre-acquired knowledge to adapt to Oulu-NPU domain using very few tasks. Interestingly, if we look at the pink rectangle (SiW dataset) we observe that training using tasks from Oulu-NPU increases significantly the performance for SiW, showing a high degree of correlation between both datasets and zero-shot properties between different domains or types of attacks (PAIs) that are somehow correlated. The same behavior appears for Replay-Mobile (tasks inside the red rectangle).

The region highlighted in purple corresponds to the CS-MAD dataset where there is only one type of PA (masks). In this scenario, CM-PAD is able to learn generalized features for mask attacks as soon as new tasks from different datasets contains similar attacks (*i.e*. 3DMAD and HKBU), leveraging previous knowledge to generalize to unseen attacks.

Regarding the influence of catastrophic forgetting in CM-PAD, we observe that the values below the main diagonal in Fig. 5 remain stable when the tasks evolve, that is, the proposed model is able to retain a high accuracy for past tasks, while being able to perform well in the current task and in unseen scenarios (values above the main diagonal)

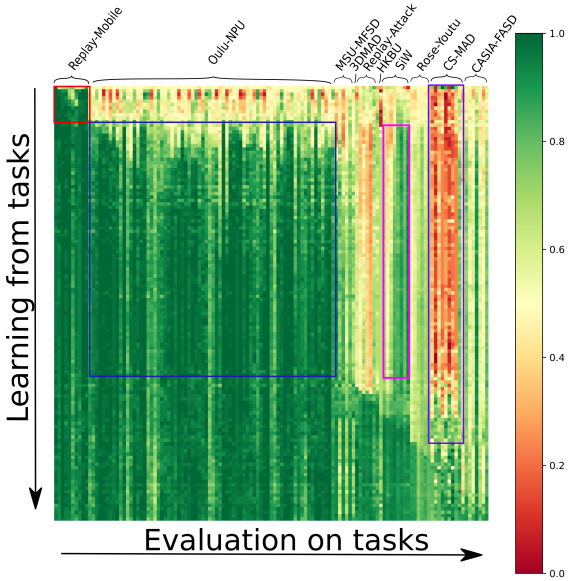

Figure 5: Confusion matrix of the proposed CM-PAD where each value corresponds to the accuracy in the Continual-Meta-Grandtest protocol of GRAD-GPAD. The tasks associated to each dataset are sorted as follows: Replay Mobile, Oulu-NPU, MSU-MFSD, 3DMAD, Replay-Attack, HKBU, SiW, Rose-Youtu, CS-MAD and CASIA-FASD (best viewed in color).

This behaviour reveals the following properties of the proposed CM-PAD: 1) catastrophic forgetting has been drastically mitigated; and 2) we are able to transfer knowledge to unseen scenarios (attacks, domains, capture devices, *etc*.), rapidly adapting to new situations.

4.3.2 Evaluation on changing PAIs

The Continual-Meta-Unseen PAI protocol. One of the most challenging scenarios for PAD approaches appears

when the models have to deal with unseen PAIs. To simulate this scenario we follow the same procedure as in the Continual-Meta-Grandtest protocol but sequentially presenting samples from different groups of PAIs. We arrange the different PAIs in three groups (*i.e*. Print, Replay and Mask) and we generate their corresponding tasks (53 for Print, 59 for Replay and 15 for Mask). Using this protocol we evaluate the influence of catastrophic forgetting on previous attack types (PAIs) when we learn a new one.

In Fig. 6 we compare our approach using the continual meta-learning setting, *Ours-CM-PAD*, with the proposed backbone using a standard fine-tuning process, *Ours-supervised* (we use disjoint splits for training/fine-tuning and testing to avoid potential overlapping identities). Note that this setting benefits the fully supervised baseline when we test on the same PAI types that it was trained on (or fine-tuned). We report ACER values for 3 sequential scenarios simulating a non i.i.d. stream of data, Figs. 6 a), b) and c), respectively. In particular, *Ours-supervised* (orange bars) is trained and evaluated as follows. In Fig. 6 a) we train and test the baseline using only Print PAIs (no catastrophic forgetting is revealed in this plot). For Fig. 6 b) we fine-tune the model obtained in the previous stage using only samples from Replay PAIs. The corresponding group of bars on the left show the performance on Print attacks after the fine-tuning process (and thus the influence of catastrophic forgetting, specially for the standard *Ours-supervised*), whereas on the right group of bars, we show the error rate while testing on Replay PAIs (same as in training). Finally, in Fig. 6 c) the first two pairs of groups of bars show the performance on Print and Replay PAIs after adapting the previous model on the new Mask PAIs, while the last pair of bars shows the performance while testing in the same PAI as in training, *i.e*. Masks.

In summary, from Fig. 6 we can see that the standard strategy, based on fine-tuning, presents a serious problem of catastrophic forgetting. It can clearly be observed that the orange bars increase dramatically each time a new task is introduced. Our model (blue bars), however, has been trained so that catastrophic forgetting is mitigated (retaining knowledge from previous PAs), while maintaining an acceptable accuracy for new tasks that may appear.

We show in Table 3 that in terms of BTI, the traditional supervised approach (*Ours-supervised*) obtains a high negative value of -18.57% while the continual meta-learning approach (*Ours-CM-PAD*) keeps a very reasonable value of -1.67%. Note that the higher the BTI is the better capacity the model has to retain past knowledge, *i.e*. less catastrophic forgetting. In other words, *Ours CM-PAD* is able to learn new knowledge losing only 1.67% average accuracy in the learned tasks, while *Ours-supervised* reduces its average performance by 18.57%. We expect that this difference in BTI would be greater if more data or new attacks were available.

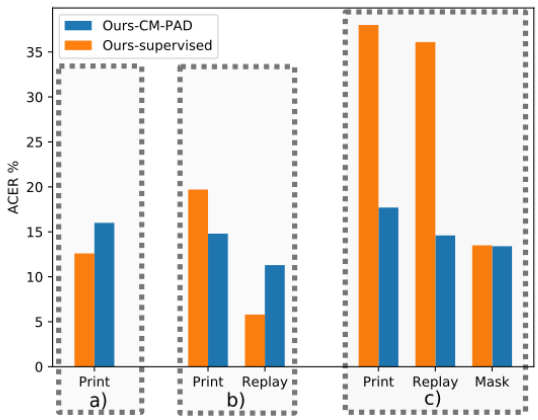

Figure 6: Evaluation on a continual learning setting. Orange bars show the performance of *Ours-supervised* (ACER) in the following sequential setting: a) trained and tested on Print PAIs, b) pretrained on Print, fine-tuned on Replay and tested on Replay PAIs and c) pretrained on Print and Replay, finetuned on Mask and tested on Mask PAIs. Blue bars show the performance (ACER) of *Ours-CM-PAD* in a continual setting where tasks from the different PAIs are presented sequentially. We use separated splits for training and testing.

	Ours-supervised	Ours-CM-PAD
BTI	-18.57	-1.67

Table 3: Backward Transfer and Interference values for the "Unseen Attack" protocol of GRAD-GPAD.

5. Conclusions

In this work we formulate for the first time the face presentation attack detection problem as a continual meta-learning task. Given the inherent dynamic nature of the problem, we introduce a new exhaustive and hierarchical categorization of all the datasets aggregated in the GRAD-GPAD framework and propose new protocols adapted to both the continual and the meta-learning settings. With a thorough experimental evaluation on GRAD-GPAD we demonstrate that our approach is not only able to perform on par to previous works on traditional protocols but also it is able to leverage past knowledge to alleviate the effects of catastrophic forgetting in continual learning settings while showing good generalization performance in related unseen scenarios. Besides, we show that we are able to rapidly adapt to new domains and attacks using a reduced set of tasks from the same scenario.

References

[1] S. M. Anjith George. Deep pixel-wise binary supervision for face presentation attack detection. In *ICB 2019*, 2019. 5

[2] L. Bertinetto, J. F. Henriques, P. Torr, and A. Vedaldi. Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*, 2019. 2

[3] S. Bhattacharjee, A. Mohammadi, and S. Marcel. Spoofing Deep Face Recognition With Custom Silicone Masks. In *BTAS*, 2018. 6

[4] Z. Boulkenafet, J. Komulainen, and A. Hadid. Face spoofing detection using colour texture analysis. *IEEE Transactions on Information Forensics and Security*, 2016. 5, 7

[5] Z. Boulkenafet, J. Komulainen, L. Li, X. Feng, and A. Hadid. OULU-NPU: A mobile face presentation attack database with real-world variations. 2017. 5, 6

[6] I. Chingovska, A. Anjos, and S. Marcel. On the effectiveness of local binary patterns in face anti-spoofing. 2012. 6

[7] A. Costa-Pazo, S. Bhattacharjee, E. Vazquez-Fernandez, and S. Marcel. The replay-mobile face presentation-attack database. In *BioSIG*, 2016. 5, 6

[8] A. Costa-Pazo, D. Jiménez-Cabello, E. Vazquez-Fernandez, J. L. Alba-Castro, and R. J. López-Sastre. Generalized presentation attack detection: a face anti-spoofing evaluation proposal. In *2019 International Conference on Biometrics*, 2019. 2, 3, 4

[9] A. Costa-Pazo, E. Vazquez-Fernandez, J. L. Alba-Castro, and D. González-Jiménez. Challenges of face presentation attack detection in real scenarios, 2019. 2

[10] T. de Freitas Pereira, A. Anjos, J. M. De Martino, and S. Marcel. Lbp-top based countermeasure against face spoofing attacks. In *Asian Conference on Computer Vision*, 2012. 2

[11] T. de Freitas Pereira, A. Anjos, J. M. D. Martino, and S. Marcel. Can face anti-spoofing countermeasures work in a real world scenario? In *ICB 2013*, 2013. 2

[12] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv preprint:1909.08383*, 2019. 3

[13] N. Erdogmus and S. Marcel. Spoofing in 2d face recognition with 3d masks and anti-spoofing with kinect. 2013. 6

[14] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2017. 2, 3

[15] J. Galbally, S. Marcel, and J. Fierrez. Biometric antispoofing methods: A survey in face recognition. *IEEE Access*, 2014. 2

[16] J. Galbally, S. Marcel, and J. Fierrez. Image quality assessment for fake biometric detection: Application to iris, fingerprint, and face recognition. *IEEE Transactions on Image Processing*, 2014. 5

[17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. 2016. 5

[18] A. Jourabloo, Y. Liu, and X. Liu. Face de-spoofing: Anti-spoofing via noise modeling. *arXiv preprint:1807.09968*, 2018. 2, 6

[19] H. Li, P. He, S. Wang, A. Rocha, X. Jiang, and A. C. Kot. Learning generalized deep feature representation for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 2018. 2

[20] H. Li, W. Li, H. Cao, S. Wang, F. Huang, and A. C. Kot. Unsupervised domain adaptation for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 2018. 2, 6

[21] S. Liu, P. C. Yuen, S. Zhang, and G. Zhao. 3d mask face anti-spoofing with remote photoplethysmography. In *Computer Vision – ECCV 2016*, 2016. 6

[22] Y. Liu*, A. Jourabloo*, and X. Liu. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In *In Proceeding of IEEE Computer Vision and Pattern Recognition*, 2018. 1, 5, 6, 7

[23] D. Lopez-Paz and M. Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017. 2

[24] A. Nichol, J. Achiam, and J. Schulman. On first-order meta-learning algorithms. *arXiv preprint:1803.02999*, 2018. 2, 3

[25] O. Nikisins, A. Mohammadi, A. Anjos, and S. Marcel. On effectiveness of anomaly detection approaches against unseen presentation attacks in face anti-spoofing. In *ICB 2018*, 2018. 2, 5, 7

[26] D. Pérez-Cabo, D. Jiménez-Cabello, A. Costa-Pazo, and R. J. López-Sastre. Deep anomaly detection for generalized face anti-spoofing. *CoRR*, 2019. 1, 2, 6, 7

[27] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, 2019. 2

[28] R. Ramachandra and C. Busch. Presentation attack detection methods for face recognition systems: A comprehensive survey. *ACM Computing Surveys*, 2017. 2

[29] M. Riemer, I. Cases, R. Ajemian, M. Liu, I. Rish, Y. Tu, and G. Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *International Conference on Learning Representations (ICLR)*, 2019. 2, 4, 5

[30] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell. Meta-learning with latent embedding optimization. In *International Conference on Learning Representations*, 2019. 2

[31] J. S. Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 1985. 4

[32] D. Wen, H. Han, and A. Jain. Face Spoof Detection with Image Distortion Analysis. *IEEE Transactions on Information Forensics and Security*, 2015. 5, 6

[33] Y. Wu and K. He. Group normalization. In *ECCV*, 2018. 5

[34] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 2016. 5

[35] Z. Zhang, J. Yan, S. Liu, Z. Lei, D. Yi, and S. Z. Li. A face antispoofing database with diverse attacks. In *ICB 2012*, 2012. 6