

Gradient Alignment for Cross-Domain Face Anti-Spoofing

Binh M. Le

Simon S. Woo*

Computer Science & Engineering Department, Sungkyunkwan University, South Korea

{bmlle, swoo}@g.skku.edu

Abstract

Recent advancements in domain generalization (DG) for face anti-spoofing (FAS) have garnered considerable attention. Traditional methods have focused on designing learning objectives and additional modules to isolate domain-specific features while retaining domain-invariant characteristics in their representations. However, such approaches often lack guarantees of consistent maintenance of domain-invariant features or the complete removal of domain-specific features. Furthermore, most prior works of DG for FAS do not ensure convergence to a local flat minimum, which has been shown to be advantageous for DG. In this paper, we introduce GAC-FAS, a novel learning objective that encourages the model to converge towards an optimal flat minimum without necessitating additional learning modules. Unlike conventional sharpness-aware minimizers, GAC-FAS identifies ascending points for each domain and regulates the generalization gradient updates at these points to align coherently with empirical risk minimization (ERM) gradient updates. This unique approach specifically guides the model to be robust against domain shifts. We demonstrate the efficacy of GAC-FAS through rigorous testing on challenging cross-domain FAS datasets, where it establishes state-of-the-art performance.

1. Introduction

With the increasing importance of security systems, face recognition technologies [9, 41, 61] have become ubiquitous in many industrial applications. However, these systems are vulnerable to presentation attacks, such as printed faces [29], and 3D masks [17], among others. Consequently, face anti-spoofing (FAS) has emerged as an essential technology to safeguard recognition systems over the past decades [6, 8, 13, 27, 37, 47, 68]. Although existing methods have achieved promising performance, they often suffer from poor generalization when exposed to unseen environments. This limitation is largely due to their assumption

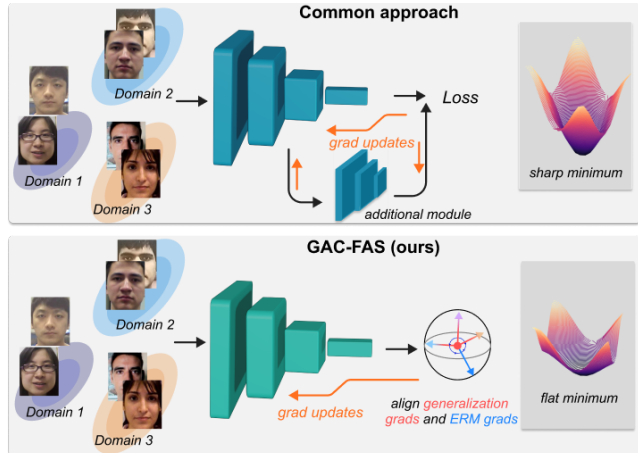


Figure 1. **Illustration of Our Learning Objective.** Most SoTA methods [43, 65, 75] for DG in FAS rely on auxiliary modules to learn domain-invariant features, and do not guarantee convergence towards a flat minimum. In contrast, our method coherently aligns the generalization gradients at ascending points of each domain with gradients derived from ERM. This approach ensures that the model converges to an optimal flat minimum and is robust against domain shifts.

of stationary settings, such as lighting conditions or sensor variations, which often do not hold in real-world scenarios.

To address this challenge, recent studies have focused on improving domain generalization (DG) for FAS by learning domain-invariant features from source training domains [5, 23, 43, 54, 56, 58, 62, 65, 75]. Primary efforts include removing domain-specific features from representations through adversarial training [23, 53, 59] or meta-learning [5, 54, 62]. Subsequent works have applied metric learning methods [56, 58] and style ensemble techniques [65, 75] to enhance robustness under domain shifts. However, most of these studies assume that domain-invariant features are preserved for DG through their specific designs of additional learning modules, without ensuring convergence of the model to a local flat minimum.

In this research, we introduce a novel training objective, namely Gradient Alignment for Cross-Domain Face Anti-Spoofing (GAC-FAS), designed to guide detectors to

*Corresponding author

wards an optimal flat minimum robust towards domain shift. This approach is particularly motivated by the recent advancements in Sharpness-Aware Minimization (SAM) [14, 30, 80], which offers a promising alternative to empirical risk minimization (ERM) [57] for seeking generalizable minima. Our objective function for DG in FAS is carefully modulated by considering the limitations of current SAM variants. When SAM is applied to entire datasets, it may produce biased updates due to the dominance of a particular domain or generate inconsistent gradients when applied to individual domains. Moreover, the updates to the SAM generalization gradient have a tendency to yield a model that is capable of handling many forms of noise, including label noise and adversarial noise. However, our primary focus is on addressing domain shifts in the context of DG for FAS.

Consequently, we propose two essential conditions for DG in FAS. First, the objective should aim for an optimal flat minimum that is both flat and low in terms of the training loss. Second, the SAM generalization gradient updates, derived at ascending points (see its definition in Sec. 3.2) for each domain, should be coherently aligned with each other and with the ERM gradient update as illustrated in Fig. 1. This dual approach enables our model to learn a more stable local minimum and become more robust to domain shifts over different face spoofing datasets.

Our comprehensive experiments on benchmark datasets under various settings, including leave-one-out, limited source domain, and performance upon convergence, demonstrate the superiority of our method compared to current state-of-the-art (SoTA) baselines. The main contributions of our work are summarized as follows¹:

1) We offer a new perspective for cross-domain FAS, shifting the focus from learning domain-invariant features to finding an optimal flat minimum for significantly improving the generalization and robustness to domain shifts.

2) We propose a novel training objective: self-regulating generalization gradient updates at ascending points to coherently align with the ERM gradient update, benefiting DG in FAS.

3) We demonstrate that our approach outperforms well-known baselines in both snapshot and convergence performance across popular FAS evaluation protocol settings.

Our paper is structured as follows: Sec. 2 provides a concise review of the most pertinent literature in FAS, along with the relationship between loss landscape sharpness and model generalization. Sec. 3 introduces the preliminaries of ERM and SAM within the context of FAS, followed by a detailed presentation of our approach. Experimental results and ablation studies are discussed in Sec. 4. Finally, we draw our conclusions in Sec. 5

¹The codes is available at: <https://github.com/leminhbinh0209/CVPR24-FAS>

2. Related Work

2.1. Face Anti-Spoofing

In the initial phases of research, handcrafted features were primarily employed as artifacts for detection. Such features include LBP [2, 8], HOG [28, 68], and SIFT [50]. Concurrently, studies have examined predefined biometric traits and behaviors, such as eye blinking [47], lip motion [27], head turning, and facial expression variations [6]. With the advent of deep neural networks, there has been a notable enhancement in detection capabilities [13, 35, 49]. Such improvements were further facilitated through diverse supervisory inputs, encompassing depth maps [69], reflection maps [72], and R-PPG signals [37]. Recently, transformer-based models have emerged, demonstrating superior efficacy in identifying spoofing attempts [21, 36].

Lately, there has been a growing interest in model generalization across disparate domains. A significant body of work has employed domain adaptation (DA) techniques, where pre-trained models are fine-tuned to novel domains using additional data [18, 32, 60, 77]. Concurrently, domain generalization (DG) methodologies, particularly those incorporating adversarial loss, have sought to achieve generalization by extracting domain-invariant features from source training domains [23, 43, 56, 58, 65, 75]. Additionally, several research works have considered meta-learning as a form of regularization to counteract domain shifts during the training phase [5, 54, 62], and others have pursued self-supervised learning to reduce reliance on labeled data [38, 44].

In contrast to prior studies in DG that primarily centered on creating auxiliary modules to eliminate domain-specific features, these approaches may not generalize effectively to unseen domains due to uncertainties in training model convergence to flat loss regions. On the other hand, our approach leverages the inherent sharpness of a model within specific domains by aligning these models. Finally, our goal is to construct a more robust and universally generalizable model for FAS.

2.2. Sharpness and Generalization

The relationship between sharpness and model generalization was initially broached in [20]. Building on this foundation and under the *i.i.d* assumption, numerous theoretical and empirical investigations delved into the relationship from the lens of loss surface geometry [4, 12, 14, 15, 22, 24]. Notably, both stochastic weight averaging (SWA) [22] and stochastic weight averaging densely (SWAD) [4] have posited, both theoretically and practically, that a flatter minimum can narrow the DG gap, leading them to propose distinct weight averaging methodologies. Nevertheless, these strategies did not explicitly encourage the model to converge towards flatter minima during its training phase.

Concurrently, Sharpness-Aware Minimization (SAM) [14] and its subsequent variants [10, 45, 63, 80] aimed to address the sharp minima problem by adjusting the objective to minimize a perturbed loss, $\mathcal{L}_p(\theta)$, which is defined as the maximum loss within a neighborhood parameter space. Specifically, Look-SAM [45] and ESAM [10] reduced the computational demands of SAM. However, they retained SAM’s primary challenge, wherein the perturbed loss $\mathcal{L}_p(\theta)$ could potentially disagree with the actual sharpness measure. To overcome this challenge, GSAM [80] minimized a surrogate gap, $h(\theta) \triangleq \mathcal{L}_p(\theta) - \mathcal{L}(\theta)$, albeit at the expense of increasing $\mathcal{L}(\theta)$. Later on, SAGM [63] rectified the inconsistencies observed in GSAM by incorporating gradient matching, ensuring model convergence to flatter regions. While SAM-based techniques have shown promise in generalizing from a single source and dealing with various types of noise, including label and adversarial noise, their application in multi-source domain DG for FAS has not been explored. Inspired by these foundational studies, we have tailored SAM to ensure that its generalization gradient updates derived from multi-source domains are aligned with each another, and with the ERM gradient update. To the best of our knowledge, this study pioneers the exploration of SAM’s capacity for DG in FAS.

3. Methods

In this section, we first define the general empirical risk for training cross-domain FAS problems. Next, we revisit variations of Sharpness-Aware Minimization for domain generalization from which we draw our motivation. Furthermore, we propose our approach, GAC-FAS, specifically tailored for the problem of DG for FAS. Finally, we analyze the benefits and prove the convergence rate of our algorithm.

3.1. Problem Definition

We begin by introducing the notion of cross-domain FAS. Consider an input space $\mathcal{X} \in \mathbb{R}^d$ and an output space $\mathcal{Y} = \{0 \text{ (fake or spoofed)}, 1 \text{ (live)}\}$. Assuming there are k distinct source domains for training, represented as $\mathcal{S} = \{\mathcal{S}_i\}_i^k$, and a singular target domain denoted by \mathcal{T} .

A neural network, characterized as $f : \mathcal{X} \rightarrow \mathcal{Y}$, is parameterized by learning parameters θ . Its aim is to distinguish whether an input x from the source domains is live or spoofed (fake). A standard approach to optimization involves the empirical risk minimization (ERM) framework [57], which aims to minimize the loss described by:

$$\min_{\theta} \mathcal{L}(\theta; \mathcal{S}) = \min_{\theta} \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}(\theta; \mathcal{S}_i), \quad (1)$$

where $\mathcal{L}(\theta; \mathcal{S}_i) = \mathbb{E}_{(x,y) \sim \mathcal{S}_i} \ell(f(x; \theta), y)$ is domain-wise empirical loss on domain i – th , and ℓ could be cross-entropy loss [23] or L_1 regression loss [16].

To minimize the empirical risk $\mathcal{L}(\theta; \mathcal{S})$, the neural network f aspires to identify the optimal parameter set θ^* . A notable challenge with ERM is its propensity to overfit the training data and converge towards sharp minima, compromising the performance on an unseen domain. Such tendencies might arise due to domain-specific attributes such as camera configurations or image resolution [56]. Consequently, targeting flatter minima when training on source domains becomes pivotal for addressing the DG for FAS problem.

3.2. Preliminaries: Sharpness-Aware Minimization

The Sharpness-Aware Minimization (SAM) [14] method aims to identify a flatter area in the vicinity of the minimum that exhibits a lower loss value. In order to attain this objective, given a training set \mathcal{D} (which can be considered as either \mathcal{S} or \mathcal{S}_i in our problem later), SAM addresses the subsequent min-max problem:

$$\min_{\theta} \mathcal{L}_p(\theta; \mathcal{D}) + \mathcal{R}(\theta; \mathcal{D}), \text{ where} \quad (2)$$

$$\mathcal{L}_p(\theta; \mathcal{D}) = \max_{\epsilon \in \mathcal{B}(\theta, \rho)} \mathcal{L}(\theta + \epsilon; \mathcal{D}), \quad (3)$$

and $\mathcal{R}(\theta; \mathcal{D})$ is a regularization term, and $\mathcal{B}(\theta, \rho) = \{\epsilon : \|\epsilon - \theta\| \leq \rho\}$ is the vicinity of model weight vector θ with a predefined constant radius ρ . Intuitively, for a given θ , the maximization in Eq. 3 identifies the most adversarial weight perturbation, denoted as ϵ^* , within the ball \mathcal{B} of radius ρ . This perturbation maximizes the empirical loss, leading to $\mathcal{L}(\theta + \epsilon^*; \mathcal{D})$ being the supremum in $\mathcal{B}(\theta, \rho)$. We now refer ϵ^* and $\theta + \epsilon^*$ as **ascending vector** and **ascending point**, respectively. By minimizing $\mathcal{L}(\theta + \epsilon^*; \mathcal{D})$, the approach encourages the selection of θ values that are situated in a region with a flatter loss landscape. Consequently, the function f exhibits enhanced stability under domain shifts, making it more resilient to unseen domains.

SAM uses Taylor expansion of the empirical loss around θ to estimate ϵ^* as follows [14]:

$$\hat{\epsilon} = \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|} \approx \arg \max_{\epsilon \in \mathcal{B}(\theta, \rho)} \mathcal{L}(\theta + \epsilon; \mathcal{D}). \quad (4)$$

Therefore, the perturbation loss of SAM reduces to

$$\mathcal{L}_p(\theta; \mathcal{D}) = \mathcal{L}(\theta + \hat{\epsilon}; \mathcal{D}), \text{ where } \hat{\epsilon} = \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{D})}{\|\nabla \mathcal{L}(\theta; \mathcal{D})\|}. \quad (5)$$

Subsequently, Zhuang *et al.* [80] and Wang *et al.* [63] introduce a surrogate gap (sharpness) $h(\theta) \triangleq \mathcal{L}_p(\theta) - \mathcal{L}(\theta)$, albeit at the expense of increasing $\mathcal{L}(\theta)$ into their objectives to obtain a better minimum. Nevertheless, there are the following limitations and issues when applying SAM and its variants for DG for FAS, as discussed subsequently.

Our Preliminary Observations and Analysis. (i) **Compensation for Input Changes:** Optimizing $\mathcal{L}_p(\theta; \mathcal{D})$ acts

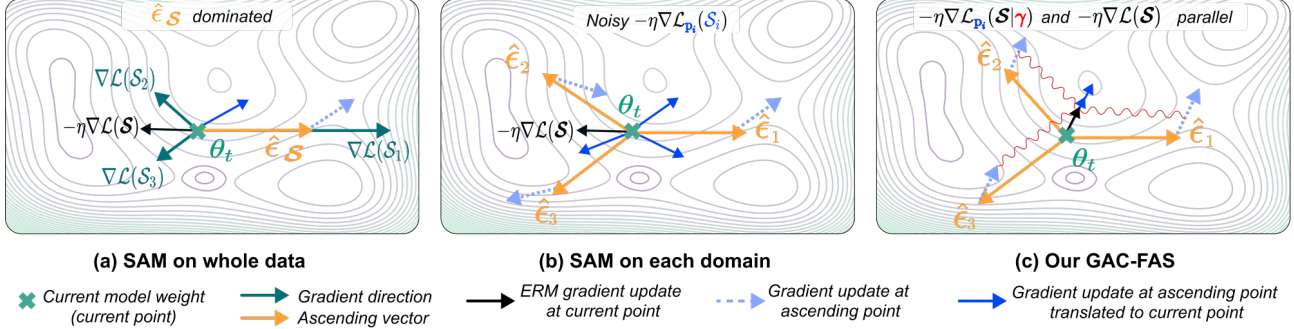


Figure 2. **Illustration of Different SAM Objective Approach.** (a) Standard SAM variants applied to the entire source dataset yield a biased ascending vector $\hat{\epsilon}_{\mathcal{S}}$, predominantly influenced by a particular domain. (b) Domain-specific gradient adjustments lead to noisy gradient estimates, impeding optimization progress. (c) Our proposed GAC-FAS addresses these issues by computing perturbation losses across the dataset at all ascending points, while concurrently adjusting gradients to align with the ERM gradients at the current point (with γ), making the model robust to domain shift.

as a compensatory mechanism for a range of input changes. These include domain shifts and various types of corruptions such as adversarial noises [67], heavy compression [31], and label noises [14]. However, our primary objective in this research is to enhance robustness against domain shifts specific to FAS. (ii) **Dominated $\hat{\epsilon}$ in All-Domain Application ($\mathcal{D} = \mathcal{S}$)**: By applying Eq. 5 or its variants to the problem of DG for FAS as formulated in Eq. 1, we can determine the optimal ascending vector $\hat{\epsilon}$ for the training source dataset $\mathcal{D} = \mathcal{S}$. However, a significant challenge arises from imbalances in dataset sizes and the presence of subtle, domain-specific artifacts. This can lead to ‘learning shortcuts,’ where the model preferentially learns from less complex or larger domains. Consequently, the ascending vector $\hat{\epsilon}$ from (θ, \mathcal{S}) can become dominated by a particular domain \mathcal{S}_i . This behaviour is similar with the long-tail problem [78, 79]. (iii) **Gradient Conflicts in Domain-wise Application ($\mathcal{D} = \mathcal{S}_i$)**: On the other hand, if we apply Eq. 5 or its variants on domain-wise manner, *i.e.*, $\mathcal{D} = \mathcal{S}_i$, the gradients of the model, derived from each domain as their ascending points, can counteract one another, leading to potential conflicts between domains. This phenomenon is further illustrated in Fig. 2. In light of these intricacies and limitations, we argue that the SAM variants might not deliver optimal generalization performance for the DG for FAS problem.

3.3. Objective of GAC-FAS

Given the insights gleaned from our prior analysis, we introduce two pivotal conditions to ensure that our model remains robust across unseen domains: (i) **Optimal minimum**: The identified minimum should not only be sufficiently low but should also reside on a flat loss surface. (ii) **Aligned cross-domain gradients**: From the source training datasets, the generalization gradient update learned from some domains should align with the ERM gradient of another domain.

Intuitively, the first condition resonates with the core objectives of optimal minimum of loss landscape from training set, as discussed previously [63, 80]. The second condition serves dual purposes. It aims to harmonize the optimization of $h(\theta)$, reducing potential conflicts. Simultaneously, it aspires for $\mathcal{L}_p(\theta; \mathcal{S})$ to exclusively compensate for domain shifts in FAS.

To fulfill these conditions, we introduce a novel optimization objective for our DA FAS, expressed as:

$$\mathcal{L}(\theta; \mathcal{S}) + \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}_{p_i}(\theta - \gamma \nabla \mathcal{L}(\theta; \mathcal{S}); \mathcal{S}) + \mathcal{R}(\theta; \mathcal{S}). \quad (6)$$

In this formulation, which is inspired by [63], we perturb the model weights *w.r.t* each individual domain. Simultaneously, we incorporate an auxiliary ERM’s gradient term, $\gamma \nabla \mathcal{L}(\theta; \mathcal{S})$, computed over all source training domains. And, Eq. 6 can be further expressed as:

$$\mathcal{L}(\theta; \mathcal{S}) + \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}(\theta + \hat{\epsilon}_i - \gamma \nabla \mathcal{L}(\theta; \mathcal{S}); \mathcal{S}) + \mathcal{R}(\theta; \mathcal{S}), \quad (7)$$

where the optimal perturbation is characterized by $\hat{\epsilon}_i = \rho \frac{\nabla \mathcal{L}(\theta; \mathcal{S}_i)}{\|\nabla \mathcal{L}(\theta; \mathcal{S}_i)\|}$ as defined in Eq. 4.

3.4. Benefits & Convergence of GAC-FAS

In this section, we offer analyses for a deeper understanding of our proposed losses for DA in FAS, detailing how they satisfy the two conditions outlined in Sec. 3.3. Subsequently, we present our novel training algorithm and the theorem regarding its convergence rate.

3.4.1 Benefits of GAC-FAS

We perform the first order Taylor expansion around $\theta + \hat{\epsilon}_i$ for the second term in Eq. 6 as follows:

$$\begin{aligned} & \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}_{p_i}(\theta - \gamma \nabla \mathcal{L}(\theta; \mathcal{S}); \mathcal{S}) \\ & \approx \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}_{p_i}(\theta; \mathcal{S}) - \gamma \langle \nabla \mathcal{L}_{p_i}(\theta; \mathcal{S}), \nabla \mathcal{L}(\theta; \mathcal{S}) \rangle. \end{aligned} \quad (8)$$




Figure 3. Illustration of the effects described by Eq. 10. The term $-\eta\nabla\mathcal{L}_{p_1}(\mathcal{S}_2)$ represents the generalization gradient update of the model learned from \mathcal{S}_1 and \mathcal{S}_2 . The term $-\eta\nabla\mathcal{L}(\mathcal{S}_3)$ denotes the ERM update for domain \mathcal{S}_3 and serves as a comparative oracle for domain shift. In the absence of our regularization, the generalization update is not robust to the domain shift associated with \mathcal{S}_3 (left), as their update directions are different. Conversely, with our regularization as formulated in Eq. 10, the generalization update aligns with the ERM update on \mathcal{S}_3 , suggesting that the model updates in a direction that is robust to domain shifts (right).

As a result, our objective in Eq. 6 can be expressed as follows:

$$\mathcal{L}(\theta; \mathcal{S}) + \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathcal{L}_{p_i}(\theta; \mathcal{S}) - \gamma \langle \nabla \mathcal{L}_{p_i}(\theta; \mathcal{S}), \nabla \mathcal{L}(\theta; \mathcal{S}) \rangle + \mathcal{R}(\theta; \mathcal{S}). \quad (9)$$

The underlying objectives of our Eq. 9 is twofold: we aim to minimize the loss functions, namely $\mathcal{L}(\theta; \mathcal{S})$ and $\mathcal{L}_{p_i}(\theta; \mathcal{S})$, while simultaneously maximizing the inner products between the gradient of $\mathcal{L}_{p_i}(\theta; \mathcal{S})$ and the gradients $\mathcal{L}(\theta; \mathcal{S})$. Specifically, condition (i) is satisfied by minimizing $\mathcal{L}(\theta; \mathcal{S})$ and $\mathcal{L}_{p_i}(\theta; \mathcal{S}), \forall i$. Our distinct contribution, however, lies in our method’s emphasis on cross-domain gradient alignment. This facet of our methodology particularly benefits the DG for FAS as elucidated subsequently. By noting that $\nabla \mathcal{L}(\theta; \mathcal{S}) = \sum_{m=1}^k \nabla \mathcal{L}(\theta; \mathcal{S}_m)$, maximizing $\langle \nabla \mathcal{L}_{p_i}(\theta; \mathcal{S}), \nabla \mathcal{L}(\theta; \mathcal{S}) \rangle$ is similar with maximizing following term:

$$\sum_{m=1}^k \sum_{n=1}^k \langle \nabla \mathcal{L}_{p_i}(\theta; \mathcal{S}_m), \nabla \mathcal{L}(\theta; \mathcal{S}_n) \rangle. \quad (10)$$

We interest in the case where \mathcal{S}_n is different with \mathcal{S}_m and \mathcal{S}_i . Maximizing the inner product $\langle \nabla \mathcal{L}_{p_i}(\theta; \mathcal{S}_m), \nabla \mathcal{L}(\theta; \mathcal{S}_n) \rangle$ implies that the *generalization gradient* update, learned from \mathcal{S}_i (via $\hat{\epsilon}_i$) and \mathcal{S}_m , must align with the *ERM gradient* of another domain, \mathcal{S}_n . This ERM gradient acts as a comparative oracle for domain shifts, guiding generalization update of the model converge towards a minimum that is robust against domain shifts. A toy example of this effect is illustrated in Fig. 3. Moreover, maximizing Eq. 10 indirectly leads to the matching each pair of $\{\nabla \mathcal{L}(\theta, \mathcal{S}_n)\}^k$, benefiting for DG [55].

3.4.2 Convergence of GAC-FAS

Theorem 1 (Proof in Supp. material). *Suppose that the loss function $\ell(\theta_t) = \ell(f(x; \theta_t), y)$ satisfies the following*

Algorithm 1 Training pipeline for GAC-FAS.

Require: DNN f parameterized by θ , training dataset $\mathcal{S} = \{\mathcal{S}_i\}_i^k$. Learning rate η . Alignment parameter γ and radius ρ . Total number of iterations T .

- 1: **for** $t \leftarrow 1$ **to** T **do**
 - 2: Sample a mini-batch $\mathcal{B} = \mathcal{B}_{\mathcal{S}_1} + \dots + \mathcal{B}_{\mathcal{S}_k}$;
 - 3: Compute grad. of reg. term $\nabla \mathcal{R}(\theta_t; \mathcal{B})$;
 - 4: #Compute grad. for 1^{st} term of Eq. 7:
 - 5: Compute the training loss gradient on each domain $\{\nabla \mathcal{L}(\theta_t; \mathcal{B}_{\mathcal{S}_i})\}_{i=1}^k$, and sum them up to obtain $\nabla \mathcal{L}(\theta_t; \mathcal{B})$;
 - 6: #Compute grad. for 2^{nd} term of Eq. 7:
 - 7: **for** domain $i \in \{1, \dots, k\}$ **do**
 - 8: $\hat{\epsilon}_i = \rho \frac{\nabla \mathcal{L}(\theta_t; \mathcal{B}_{\mathcal{S}_i})}{\|\nabla \mathcal{L}(\theta_t; \mathcal{B}_{\mathcal{S}_i})\|}$ #ascending vector
 - 9: $\nabla \mathcal{L}_p^i = \nabla \mathcal{L}(\theta_t + \hat{\epsilon}_i - \gamma \nabla \mathcal{L}(\theta_t; \mathcal{B}); \mathcal{B})$
 - 10: **end for**
 - 11: #Update weights:
 - 12: $\theta_{t+1} = \theta_t - \eta \cdot (\nabla \mathcal{L}(\theta_t; \mathcal{B}) + \frac{1}{k} \sum_{i=1}^k \nabla \mathcal{L}_p^i + \nabla \mathcal{R}(\theta_t; \mathcal{B}))$
 - 13: $t = t + 1$
 - 14: **end for**
-

assumptions: (i) its gradient $g(\theta_t) = \nabla \ell(\theta_t)$ is bounded, i.e., $\|g(\theta_t)\| \leq G, \forall t$. (ii) The stochastic gradient is L -Lipchitz, i.e., $\|g(\theta_t) - g(\theta'_t)\| \leq L\|\theta_t - \theta'_t\|, \forall \theta_t, \theta'_t$. Let the learning rate η_t be $\frac{\eta_0}{\sqrt{t}}$, and let the perturbation be proportional to the learning rate, i.e., $\rho_t = \frac{\rho}{\sqrt{t}}$, and $\gamma_t = \frac{\gamma}{\sqrt{t}}$, we have:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathbb{E}_{(x,y) \sim \mathcal{S}_i} [\|\nabla \ell(\theta_t)\|^2] \leq \mathcal{O}\left(\frac{\log T}{\sqrt{T}}\right), \text{ and}$$

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathcal{S}_i \sim \mathcal{S}} \mathbb{E}_{(x,y) \sim \mathcal{S}_i} [\|\nabla \ell(\theta_t^{\text{adv}})\|^2] \leq \mathcal{O}\left(\frac{\log T}{\sqrt{T}}\right),$$

where $\theta_t^{\text{adv}} = \theta_t + \hat{\epsilon}_t - \gamma_t \delta_t$, $\delta_t = \sum_{j=1}^k \nabla \ell(f(x'_j; \theta_t), y'_j)$, and $(x'_j, y'_j) \sim \mathcal{S}_j$.

Theorem 1 implies that both ℓ and ℓ_p converge at rate $\mathcal{O}(\log T / \sqrt{T})$, and it matches convergence rate of first-order gradient optimizers such as Adam [26].

Equipped with Theorem 1, we present the overall training pipeline of our GAC-FAS in Algorithm 1.

4. Experimental Results

In this section, we compare our method with previous SoTA baselines using standard FAS evaluation protocol settings. Additionally, we assess the effectiveness of our algorithm in scenarios where the training source is limited. We then compare its convergence performance with other baselines. Finally, we conduct several ablation studies to explore alternatives to the minimizer for DG in FAS and the effects of hyperparameter tuning.

Methods	ICM \rightarrow O		OCM \rightarrow I		OCI \rightarrow M		OMI \rightarrow C	
	HTER \downarrow	AUC \uparrow	HTER \downarrow	AUC \uparrow	HTER \downarrow	AUC \uparrow	HTER \downarrow	AUC \uparrow
MMD-AAE [33]	40.98	63.08	31.58	75.18	27.08	83.19	44.59	58.29
MADDG [53]	27.98	80.02	22.19	84.99	17.69	88.06	24.50	84.51
RFM [54]	16.45	91.16	17.30	90.48	13.89	93.98	20.27	88.16
SSDG-M [23]	25.17	81.83	18.21	94.61	16.67	90.47	23.11	85.45
SSDG-R [23]	15.61	91.54	11.71	96.59	7.38	97.17	10.44	95.94
D ² AM [5]	15.27	90.87	15.43	91.22	12.70	95.66	20.98	85.58
SDA [62]	23.10	84.30	15.60	90.10	15.40	91.80	24.50	84.40
DRDG [40]	15.63	91.75	15.56	91.79	12.43	95.81	19.05	88.79
ANRL [39]	15.67	91.90	16.03	91.04	10.83	96.75	17.85	89.26
SSAN [65]	13.72	93.63	8.88	96.79	6.67	98.75	10.00	96.67
AMEL [76]	11.31	93.96	18.60	88.79	10.23	96.62	11.88	94.39
EBDG [11]	15.66	92.02	18.69	92.28	9.56	97.17	18.34	90.01
PathNet [58]	11.82	95.07	13.40	95.67	7.10	98.46	11.33	94.58
IADG [75]	8.86	97.14	10.62	94.50	5.41	98.19	8.70	96.40
SA-FAS [56]	10.00	96.23	6.58	97.54	5.95	96.55	8.78	95.37
UDG-FAS [43]	10.97	95.36	5.86	98.62	5.95	98.47	9.82	96.76
GAC-FAS (ours)	8.60 ^{0.28}	97.16 ^{0.40}	4.29 ^{0.83}	98.87 ^{0.60}	5.00 ^{0.00}	97.56 ^{0.06}	8.20 ^{0.43}	95.16 ^{0.09}

Table 1. Evaluation of cross-domain face anti-spoofing on four leading benchmark datasets: CASIA (C), Idiap Replay (I), MSU-MFSD (M), and Oulu-NPU (O). Methods are benchmarked for optimal performance using the standard evaluation procedure outlined in [23]. Symbols \uparrow and \downarrow signify that larger and smaller values are preferable, respectively.

Methods	MI \rightarrow C		MI \rightarrow O	
	HTER \downarrow	AUC \uparrow	HTER \downarrow	AUC \uparrow
MSLBP [46]	51.16	52.09	43.63	58.07
Color Texture [2]	55.17	46.89	53.31	45.16
LBPTOP [8]	45.27	54.88	47.26	50.21
MADDG [53]	41.02	64.33	39.35	65.10
SSDG-M [23]	31.89	71.29	36.01	66.88
D ² AM [5]	32.65	72.04	27.70	75.36
DRDG [40]	31.28	71.50	33.35	69.14
ANRL [39]	31.06	72.12	30.73	74.10
SSAN [65]	30.00	76.20	29.44	76.62
EBDG [11]	27.97	75.84	25.94	78.28
AMEL [76]	24.52	82.12	19.68	87.01
IAGD [75]	24.07	85.13	18.47	90.49
GAC-FAS (ours)	16.91 ^{1.17}	88.12 ^{0.58}	17.88 ^{0.15}	89.67 ^{0.39}

Table 2. Evaluation on limited source domains. Baseline results are sourced from [75].

4.1. Experiment Settings

Datasets. Our experiments are conducted on four benchmark datasets: Idiap Replay Attack [7] (I), OULU-NPU [3] (O), CASIA-MFSD [74] (C), and MSU-MFSD [66] (M). Consistent with prior works, we treat each dataset as a separate domain and employ a leave-one-out testing protocol to evaluate cross-domain generalization capabilities. For instance, the protocol ICM \rightarrow O involves training on Idiap Replay Attack, CASIA-MFSD, and MSU-MFSD, and testing on OULU-NPU. **Implementation Details.** Input images are detected and cropped using MTCNN [71], then resized to 256×256 pixels. We employ a ResNet-18 [19] architecture, pre-trained on the ImageNet dataset [51], as our feature extraction backbone to main-

Methods	AUC \uparrow	Method	AUC \uparrow
SVM1+IMQ [1]	70.23 ^{12.69}	Saha <i>et al.</i> [52]	79.20
CDCN [70]	88.69 ^{10.56}	Panwar <i>et al.</i> [48]	80.00
CDCN++ [70]	87.53 ^{10.90}	SSDG-R [23]	82.11
SSAN [65]	88.01 ^{9.93}	CIFAS [42]	83.20
TTN-S [64]	89.71 ^{9.17}	UDG-FAS [43]	87.26
UDG-FAS [43]	92.43 ^{6.86}	GAC-FAS (ours)	89.27 ^{0.58}
GAC-FAS (ours)	93.39 ^{4.27}		

(a) Unseen 2D attack

(b) Unseen 3D attack

Table 3. Evaluation on unseen attacks. Baseline results are sourced from [43].

tain consistency with SoTA baselines [23, 56, 65, 75]. The network is trained using an SGD optimizer with an initial learning rate of 0.005. Our regularization strategy includes weight decay and supervised contrastive learning, applied at an intermediate layer, to promote inter-domain discriminability [25, 56]. The hyperparameters are set as follows: $\{\gamma = 0.0002, \rho = 0.1\}$. We run each experiment three times and take the average performance to report.

Evaluation Metrics. Model performance is quantified using three standard metrics: Half Total Error Rate (HTER), Area Under the Receiver Operating Characteristic Curve (AUC), and True Positive Rate (TPR) at a False Positive Rate (FPR) of 5%, denoted as TPR95.

4.2. Comparison to SoTA baselines

Leave-One-Out. Table 1 presents a comprehensive comparison with a broad range of recent studies addressing DG in FAS. Based on the results, we make the following observations: (1) Recent advancements in FAS methods have

Methods	ICM \rightarrow O			OCM \rightarrow I			OCI \rightarrow M			OMI \rightarrow C		
	HTER \downarrow / AUC \uparrow / TPR95 \uparrow	HTER \downarrow / AUC \uparrow / TPR95 \uparrow	HTER \downarrow / AUC \uparrow / TPR95 \uparrow	HTER \downarrow / AUC \uparrow / TPR95 \uparrow	HTER \downarrow / AUC \uparrow / TPR95 \uparrow	HTER \downarrow / AUC \uparrow / TPR95 \uparrow						
SSDG-R [23]	15.83 ^{1.29} / 92.13 ^{0.96} / 66.54 ^{4.00}	14.65 ^{1.21} / 91.93 ^{1.35} / 53.68 ^{2.56}	22.84 ^{1.14} / 78.67 ^{1.31} / 50.80 ^{5.95}	28.76 ^{0.89} / 80.91 ^{1.10} / 41.47 ^{2.68}								
SSAN-R [65]	25.72 ^{3.74} / 79.37 ^{4.69} / 36.75 ^{5.19}	35.39 ^{8.04} / 70.13 ^{9.03} / 64.00 ^{2.70}	21.79 ^{3.68} / 84.06 ^{3.78} / 51.91 ^{4.28}	26.44 ^{2.91} / 78.84 ^{2.83} / 45.36 ^{4.29}								
PatchNet [58]	23.49 ^{1.80} / 84.62 ^{1.92} / 39.39 ^{6.83}	29.75 ^{2.76} / 80.53 ^{1.35} / 54.25 ^{2.18}	25.92 ^{1.13} / 83.43 ^{0.87} / 38.75 ^{8.31}	36.26 ^{1.98} / 71.38 ^{1.89} / 19.22 ^{3.85}								
SA-FAS [56]	11.29 ^{0.32} / 95.23 ^{0.24} / 73.38 ^{1.64}	11.48 ^{1.10} / 95.74 ^{0.55} / 77.05 ^{3.26}	14.36 ^{1.10} / 92.06 ^{0.53} / 55.71 ^{4.82}	19.40 ^{0.66} / 88.69 ^{0.67} / 50.53 ^{3.60}								
GAC-FAS (ours)	9.89 ^{0.47} / 96.44 ^{0.18} / 80.47 ^{1.34}	12.51 ^{3.03} / 93.03 ^{2.24} / 77.38 ^{8.50}	12.29 ^{1.29} / 95.35 ^{0.57} / 72.00 ^{3.84}	15.37 ^{1.52} / 91.67 ^{1.67} / 58.67 ^{10.55}								

Table 4. Evaluation at convergence. A comprehensive assessment of cross-domain face anti-spoofing on prominent databases: CASIA (C), Idiap Replay (I), MSU-MFSD (M), and Oulu-NPU (O). Methods are benchmarked using their mean and standard deviation performance over the final 10 evaluations. Baseline results are sourced from [56].

Model / loss	DGrad	HTER \downarrow	AUC \uparrow	TPR95 \uparrow
Fish [55]	✓	33.83 ^{0.74}	72.34 ^{0.37}	17.50 ^{1.48}
SAM [14]	✗	11.99 ^{0.49}	95.13 ^{0.16}	73.31 ^{1.57}
	✓	12.51 ^{0.89}	95.45 ^{0.35}	72.94 ^{2.99}
ASAM [30]	✗	11.73 ^{0.40}	95.36 ^{0.28}	75.36 ^{1.61}
	✓	12.18 ^{0.93}	95.40 ^{0.49}	75.92 ^{1.54}
SAGM [63]	✗	11.63 ^{0.53}	95.33 ^{0.23}	74.81 ^{2.07}
	✓	12.19 ^{0.50}	95.06 ^{0.21}	74.17 ^{1.48}
LookSAM [45]	✗	11.56 ^{0.54}	95.56 ^{0.15}	75.64 ^{1.19}
	✓	11.06 ^{0.50}	95.76 ^{0.15}	75.86 ^{1.75}
GSAM [80]	✗	11.60 ^{0.38}	95.30 ^{0.18}	74.44 ^{1.79}
	✓	12.48 ^{0.83}	95.54 ^{0.35}	73.39 ^{2.92}
Reg. $\langle \nabla \mathcal{L}_{p_i}(S_i), \nabla \mathcal{L}(S) \rangle$	✓	11.35 ^{0.54}	95.55 ^{0.18}	73.58 ^{1.21}
GAC-FAS (ours)	✓	9.89 ^{0.47}	96.44 ^{0.18}	80.47 ^{1.34}

Table 5. **Ablation study:** Alternative approach using domain gradient for DG in FAS. DGrad indicates whether domain-wise gradients are used in each method (see Fig. 2).

achieved significant breakthroughs in performance. However, a performance plateau is evident among these methods, largely because they do not incorporate the sharpness of the loss landscape into their objectives. (2) Our method consistently outperforms the majority of the surveyed DG for FAS approaches [23, 43, 56, 58, 65, 75] across all four experimental setups. Notably, we report a 1.56% improvement in the HTER (reducing from 5.86% to 4.29%) in the OCM \rightarrow I experiment, translating to an enhancement of over 26%.

Limited Source Domains. Table 2 summarizes the performance of our method when source domains are extremely limited. Adhering to the standard settings established by prior research [65, 75, 76], we utilize MSU-MFSD and Idiap Replay Attack as source domains for training, while OULU-NPU and CASIA-MFSD serve as the test datasets. Our proposed method consistently surpasses SoTA baselines, achieving substantial margins of improvement on the HTER metric. A particularly significant enhancement is observed on the MI \rightarrow C setup, where our approach yields approximately a 7% reduction in HTER (decreasing from 24.07% to 16.91%).

Unseen Attacks. In this experiment, we evaluate the detector’s performance against unseen 2D and 3D attacks.




Figure 4. **Ablation study:** Sensitivity analysis of hyper-parameters γ and ρ on ICM \rightarrow O upon convergence performance.

For 2D attacks, we adopt the ‘leave-one-attack-type-out’ method from [1], training on two domains of I, C, M, and testing on an unseen attack of unseen domain. In 3D attacks, we train on O, C, M, and evaluate using a 3D attack subset from the CelebA-Spoof dataset [73]. Results, shown in Table 3, indicate our approach outperforms baselines by significant margins of 0.96% and 2.01% on AUC metric for 2D and 3D attacks, respectively.

Comparison Upon Convergence. In their recent study, Sun *et al.* highlight that a snapshot performance report of a test set may not accurately reflect the true generalization ability of a detection model [56]. In alignment with their methodology, we report the average performance of our model across the last 10 evaluations in Table 4. As evident from the results, our method consistently demonstrates superior convergence on the three metrics: HTER, AUC, and TPR95, and achieves comparable results to SA-FAS in the OCM \rightarrow I experiment. Notably, we observe a 4% reduction in HTER in the OMI \rightarrow C experiment (from 19.40% down to 15.37%). These results suggest that our proposed GAC-FAS enables the model to converge to flatter and more stable minima compared to other approaches.

4.3. Ablation Studies

Alternatives of Optimizer. In Table 5, we explore various alternative objectives for DG in FAS, including SAM [14], ASAM [30], GSAM [80], [63], and LookSAM [45]. Furthermore, we incorporate Fish [55], a gradient matching method for DG, into our study. We experiment with ICM \rightarrow O task and report their performances upon convergence. The checkmarks in the second column denote the utilization




Figure 5. **Loss Landscape Visualization** [34]. SoTA baselines exhibit convergence to sharp minima (Figures **a** to **d**) on the training set of $\text{ICM} \rightarrow \text{O}$. In contrast, our GAC-FAS (Figure **e**) achieves convergence to a flatter minimum, indicative of potentially better generalization.

of domain-wise gradients (see Fig. 2). While most methods demonstrate improved convergence with whole-data gradients per iteration compared to domain-wise gradients, they do not unequivocally surpass the existing SoTA baselines in DG for FAS. Domain-wise gradients, as we observed, are prone to noisy gradients at ascending points, whereas whole-data gradients tend to be dominated by specific domains, which can impede convergence. Notably, Fish [55] exhibits competitive performance on other datasets but shows slower convergence on FAS datasets compared to SAM-based objectives.

Effects of Hyper-parameters γ and ρ . We investigate the sensitivities of our GAC-FAS with respect to γ and ρ , summarizing the results of our analysis in Fig. 4. This study focuses on the $\text{ICM} \rightarrow \text{O}$ task, reporting the HTER metric performance upon convergence. We test a range of values for the hyperparameters $\gamma \in \{0.0, 0.0001, 0.0002, 0.001, 0.002\}$ and $\rho \in \{0.005, 0.05, 0.1, 0.2, 0.4\}$, noting that at $\gamma = 0.0$, no regularization is applied. The results indicate that with γ set to 0.0, our model’s performance is competitive with the current SoTA as shown in Table 4, yet it does not achieve the best result. Furthermore, when the settings for γ and ρ are not excessively large, the model’s performance tends to be stable and can exceed that of previous SoTA methods. However, higher values of these hyperparameters may deteriorate performance. It is important to note that we did not fine-tune these hyperparameters to optimize test accuracy for each experimental task; thus, the same hyper-parameter settings used in Sec. 4.2 may not be optimal, even though they outperform all current SoTA methods on the datasets.

Loss Landscape Visualization. Figure 5 presents the loss landscape visualization [34] of GAC-FAS, in comparison with four SoTA approaches: SSAN [65], SSDG [43], SA-FAS [56], and IADG [75]. We employ negative log-likelihood as the loss metric and use the training set of the $\text{ICM} \rightarrow \text{O}$ task for visualization. While all baseline methods demonstrate comparable generalization capabilities to our method in Table 1, they distinctly exhibit sharp minima in

their loss functions, characterized by steep gradients in the loss landscape as shown in Fig. 5. **a**-**d**). In contrast, our proposed approach reveals a flatter minimum, which may correlate with enhanced generalization. These observations provide additional insights into the superior numerical results achieved by our method in various experiments, as discussed in Sec. 4.2.

5. Conclusion

In this paper, we introduced GAC-FAS, a novel framework designed to optimize the minimum for Domain Generalization (DG) in Face Anti-Spoofing (FAS). Inspired by recent advancements in DG methods that leverage loss sharpness-aware objectives, our approach involves identifying and utilizing ascending points for each domain within the training dataset. A key underlying novelty in our methodology is the regulation of SAM generalization gradients of whole data at these ascending points, aligning them coherently with gradients derived from ERM. Through comprehensive analysis and a series of rigorous experiments, we have demonstrated that GAC-FAS not only achieves superior generalization capabilities in FAS tasks but also consistently outperforms current SoTA baselines by significant margins. This performance consistency is observed across a variety of common experimental setups, underscoring the robustness and effectiveness of our proposed method.

Acknowledgements. This work was partly supported by Institute for Information & communication Technology Planning & evaluation (IITP) grants funded by the Korean government MSIT: (No. 2022-0-01199, Graduate School of Convergence Security at Sungkyunkwan University), (No. 2022-0-01045, Self-directed Multi-Modal Intelligence for solving unknown, open domain problems), (No. 2022-0-00688, AI Platform to Fully Adapt and Reflect Privacy-Policy Changes), (No. 2021-0-02068, Artificial Intelligence Innovation Hub), (No. 2019-0-00421, AI Graduate School Support Program at Sungkyunkwan University), and (No. RS-2023-00230337, Advanced and Proactive AI Platform Research and Development Against Malicious Deepfakes).

References

- [1] Shervin Rahimzadeh Arashloo, Josef Kittler, and William Christmas. An anomaly detection approach to face spoofing detection: A new formulation and evaluation protocol. *IEEE access*, 5:13868–13882, 2017. 6, 7
- [2] Zinelabidine Boulkenafet, Jukka Komulainen, and Abdenour Hadid. Face anti-spoofing based on color texture analysis. In *2015 IEEE international conference on image processing (ICIP)*, pages 2636–2640. IEEE, 2015. 2, 6
- [3] Zinelabinde Boulkenafet, Jukka Komulainen, Lei Li, Xiaoyi Feng, and Abdenour Hadid. Oulu-npu: A mobile face presentation attack database with real-world variations. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*, pages 612–618. IEEE, 2017. 6
- [4] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021. 2
- [5] Zhihong Chen, Taiping Yao, Kekai Sheng, Shouhong Ding, Ying Tai, Jilin Li, Feiyue Huang, and Xinyu Jin. Generalizable representation learning for mixture domain face anti-spoofing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 1132–1139, 2021. 1, 2, 6
- [6] Girija Chetty. Biometric liveness checking using multimodal fuzzy fusion. In *International Conference on Fuzzy Systems*, pages 1–8. IEEE, 2010. 1, 2
- [7] Ivana Chingovska, André Anjos, and Sébastien Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *2012 BIOSIG-proceedings of the international conference of biometrics special interest group (BIOSIG)*, pages 1–7. IEEE, 2012. 6
- [8] Tiago de Freitas Pereira, André Anjos, José Mario De Martino, and Sébastien Marcel. Lbp- top based countermeasure against face spoofing attacks. In *Computer Vision-ACCV 2012 Workshops: ACCV 2012 International Workshops, Daejeon, Korea, November 5-6, 2012, Revised Selected Papers, Part I 11*, pages 121–132. Springer, 2013. 1, 2, 6
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 1
- [10] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent Tan. Efficient sharpness-aware minimization for improved training of neural networks. In *International Conference on Learning Representations*, 2021. 3
- [11] Zhekai Du, Jingjing Li, Lin Zuo, Lei Zhu, and Ke Lu. Energy-based domain generalization for face anti-spoofing. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1749–1757, 2022. 6
- [12] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *Uncertainty in Artificial Intelligence*, 2017. 2
- [13] Litong Feng, Lai-Man Po, Yuming Li, Xuyuan Xu, Fang Yuan, Terence Chun-Ho Cheung, and Kwok-Wai Cheung. Integration of image quality and motion cues for face anti-spoofing: A neural network approach. *Journal of Visual Communication and Image Representation*, 38:451–460, 2016. 1, 2
- [14] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2020. 2, 3, 4, 7
- [15] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018. 2
- [16] Anjith George and Sébastien Marcel. Deep pixel-wise binary supervision for face presentation attack detection. In *2019 International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2019. 3
- [17] Andy Greenberg. Hackers say they’ve broken face id a week after iphone x release. <https://www.wired.com/story/hackers-say-broke-face-id-security/>, November 2017. Accessed: 2023-07-01. 1
- [18] Xiao Guo, Yaojie Liu, Anil Jain, and Xiaoming Liu. Multi-domain learning for updating face anti-spoofing models. In *European Conference on Computer Vision*, pages 230–249. Springer, 2022. 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. *Advances in neural information processing systems*, 7, 1994. 2
- [21] Hsin-Ping Huang, Deqing Sun, Yaojie Liu, Wen-Sheng Chu, Taihong Xiao, Jinwei Yuan, Hartwig Adam, and Ming-Hsuan Yang. Adaptive transformers for robust few-shot cross-domain face anti-spoofing. In *European Conference on Computer Vision*, pages 37–54. Springer, 2022. 2
- [22] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, pages 876–885. Association For Uncertainty in Artificial Intelligence (AUAI), 2018. 2
- [23] Yunpei Jia, Jie Zhang, Shiguang Shan, and Xilin Chen. Single-side domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8484–8493, 2020. 1, 2, 3, 6, 7
- [24] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2016. 2
- [25] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and

- Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 6
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [27] Klaus Kollreider, Hartwig Fronthaler, Maycel Isaac Faraj, and Josef Bigun. Real-time face detection and motion analysis with application in “liveness” assessment. *IEEE Transactions on Information Forensics and Security*, 2(3):548–558, 2007. 1, 2
- [28] Jukka Komulainen, Abdenour Hadid, and Matti Pietikäinen. Context based face anti-spoofing. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE, 2013. 2
- [29] Paul Kunert. Phones’ facial recog tech ‘fooled’ by low-res 2d photo. https://www.theregister.com/2023/05/19/2d_photograph_facial_recog/, May 2023. Accessed: 2023-07-01. 1
- [30] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, pages 5905–5914. PMLR, 2021. 2, 7
- [31] Binh M Le and Simon S Woo. Quality-agnostic deepfake detection with intra-model collaborative learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22378–22389, 2023. 4
- [32] Haoliang Li, Wen Li, Hong Cao, Shiqi Wang, Feiyue Huang, and Alex C Kot. Unsupervised domain adaptation for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 13(7):1794–1809, 2018. 2
- [33] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 6
- [34] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018. 8
- [35] Lei Li, Xiaoyi Feng, Zinelabidine Boulkenafet, Zhaoqiang Xia, Mingming Li, and Abdenour Hadid. An original face anti-spoofing approach using partial convolutional neural network. In *2016 Sixth international conference on image processing theory, tools and applications (IPTA)*, pages 1–6. IEEE, 2016. 2
- [36] Chen-Hao Liao, Wen-Cheng Chen, Hsuan-Tung Liu, Yi-Ren Yeh, Min-Chun Hu, and Chu-Song Chen. Domain invariant vision transformer learning for face anti-spoofing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6098–6107, 2023. 2
- [37] Bofan Lin, Xiaobai Li, Zitong Yu, and Guoying Zhao. Face liveness detection by rppg features and contextual patch-based cnn. In *Proceedings of the 2019 3rd international conference on biometric engineering and applications*, pages 61–68, 2019. 1, 2
- [38] Haozhe Liu, Zhe Kong, Raghavendra Ramachandra, Feng Liu, Linlin Shen, and Christoph Busch. Taming self-supervised learning for presentation attack detection: In-image de-folding and out-of-image de-mixing. *arXiv preprint arXiv:2109.04100*, 2021. 2
- [39] Shubao Liu, Ke-Yue Zhang, Taiping Yao, Mingwei Bi, Shouhong Ding, Jilin Li, Feiyue Huang, and Lizhuang Ma. Adaptive normalized representation learning for generalizable face anti-spoofing. In *Proceedings of the 29th ACM international conference on multimedia*, pages 1469–1477, 2021. 6
- [40] Shubao Liu, Ke-Yue Zhang, Taiping Yao, Kekai Sheng, Shouhong Ding, Ying Tai, Jilin Li, Yuan Xie, and Lizhuang Ma. Dual reweighting domain generalization for face presentation attack detection. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 867–873, 2021. 6
- [41] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 1
- [42] Yuchen Liu, Yabo Chen, Wenrui Dai, Chenglin Li, Junni Zou, and Hongkai Xiong. Causal intervention for generalizable face anti-spoofing. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 01–06. IEEE, 2022. 6
- [43] Yuchen Liu, Yabo Chen, Mengran Gou, Chun-Ting Huang, Yaoming Wang, Wenrui Dai, and Hongkai Xiong. Towards unsupervised domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20654–20664, 2023. 1, 2, 6, 7, 8
- [44] Yuchen Liu, Yabo Chen, Mengran Gou, Chun-Ting Huang, Yaoming Wang, Wenrui Dai, and Hongkai Xiong. Towards unsupervised domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20654–20664, 2023. 2
- [45] Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable sharpness-aware minimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12360–12370, 2022. 3, 7
- [46] Jukka Määttä, Abdenour Hadid, and Matti Pietikäinen. Face spoofing detection from single images using micro-texture analysis. In *2011 international joint conference on Biometrics (IJCB)*, pages 1–7. IEEE, 2011. 6
- [47] Gang Pan, Lin Sun, Zhaohui Wu, and Shihong Lao. Eyeblick-based anti-spoofing in face recognition from a generic webcam. In *2007 IEEE 11th international conference on computer vision*, pages 1–8. IEEE, 2007. 1, 2
- [48] Ankush Panwar, Pratyush Singh, Suman Saha, Danda Pani Paudel, and Luc Van Gool. Unsupervised compound domain adaptation for face anti-spoofing. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–8. IEEE, 2021. 6
- [49] Keyurkumar Patel, Hu Han, and Anil K Jain. Cross-database face antispoofing with robust feature representation. In *Biometric Recognition: 11th Chinese Conference, CCBR 2016, Chengdu, China, October 14-16, 2016, Proceedings 11*, pages 611–619. Springer, 2016. 2

- [50] Keyurkumar Patel, Hu Han, and Anil K Jain. Secure face unlock: Spoof detection on smartphones. *IEEE transactions on information forensics and security*, 11(10):2268–2283, 2016. [1](#), [2](#)
- [51] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [6](#)
- [52] Suman Saha, Wenhao Xu, Menelaos Kanakis, Stamatios Georgoulis, Yuhua Chen, Danda Pani Paudel, and Luc Van Gool. Domain agnostic feature learning for image and video based face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 802–803, 2020. [6](#)
- [53] Rui Shao, Xiangyuan Lan, Jiawei Li, and Pong C Yuen. Multi-adversarial discriminative deep domain generalization for face presentation attack detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10023–10031, 2019. [1](#), [6](#)
- [54] Rui Shao, Xiangyuan Lan, and Pong C Yuen. Regularized fine-grained meta face anti-spoofing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11974–11981, 2020. [1](#), [2](#), [6](#)
- [55] Yuge Shi, Jeffrey Seely, Philip Torr, N Siddharth, Awani Hanun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. In *International Conference on Learning Representations*, 2021. [5](#), [7](#), [8](#)
- [56] Yiyun Sun, Yaojie Liu, Xiaoming Liu, Yixuan Li, and Wen-Sheng Chu. Rethinking domain generalization for face anti-spoofing: Separability and alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24563–24574, 2023. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [57] Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991. [2](#), [3](#)
- [58] Chien-Yi Wang, Yu-Ding Lu, Shang-Ta Yang, and Shang-Hong Lai. Patchnet: A simple face anti-spoofing framework via fine-grained patch recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20281–20290, 2022. [1](#), [2](#), [6](#), [7](#)
- [59] Guoqing Wang, Hu Han, Shiguang Shan, and Xilin Chen. Improving cross-database face presentation attack detection via adversarial domain adaptation. In *2019 International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2019. [1](#)
- [60] Guoqing Wang, Hu Han, Shiguang Shan, and Xilin Chen. Unsupervised adversarial domain adaptation for cross-domain face presentation attack detection. *IEEE Transactions on Information Forensics and Security*, 16:56–69, 2020. [2](#)
- [61] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018. [1](#)
- [62] Jingjing Wang, Jingyi Zhang, Ying Bian, Youyi Cai, Chunmao Wang, and Shiliang Pu. Self-domain adaptation for face anti-spoofing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 2746–2754, 2021. [1](#), [2](#), [6](#)
- [63] Pengfei Wang, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. Sharpness-aware gradient matching for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3769–3778, 2023. [3](#), [4](#), [7](#)
- [64] Zhuo Wang, Qiangchang Wang, Weihong Deng, and Guodong Guo. Learning multi-granularity temporal characteristics for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 17:1254–1269, 2022. [6](#)
- [65] Zhuo Wang, Zezheng Wang, Zitong Yu, Weihong Deng, Jiahong Li, Tingting Gao, and Zhongyuan Wang. Domain generalization via shuffled style assembly for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4123–4133, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [66] Di Wen, Hu Han, and Anil K Jain. Face spoof detection with image distortion analysis. *IEEE Transactions on Information Forensics and Security*, 10(4):746–761, 2015. [6](#)
- [67] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33:2958–2969, 2020. [4](#)
- [68] Jianwei Yang, Zhen Lei, Shengcai Liao, and Stan Z Li. Face liveness detection with component dependent descriptor. In *2013 International Conference on Biometrics (ICB)*, pages 1–6. IEEE, 2013. [1](#), [2](#)
- [69] Zitong Yu, Xiaobai Li, Jingang Shi, Zhaoqiang Xia, and Guoying Zhao. Revisiting pixel-wise supervision for face anti-spoofing. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(3):285–295, 2021. [2](#)
- [70] Zitong Yu, Chenxu Zhao, Zezheng Wang, Yunxiao Qin, Zhuo Su, Xiaobai Li, Feng Zhou, and Guoying Zhao. Searching central difference convolutional networks for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5295–5305, 2020. [6](#)
- [71] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters*, 23(10):1499–1503, 2016. [6](#)
- [72] Ke-Yue Zhang, Taiping Yao, Jian Zhang, Shice Liu, Bangjie Yin, Shouhong Ding, and Jilin Li. Structure destruction and content combination for face anti-spoofing. In *2021 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–6. IEEE, 2021. [2](#)
- [73] Yuanhan Zhang, Zhenfei Yin, Yidong Li, Guojun Yin, Junjie Yan, Jing Shao, and Ziwei Liu. Celeba-spoof: Large-scale face anti-spoofing dataset with rich annotations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 70–85. Springer, 2020. [7](#)
- [74] Zhiwei Zhang, Junjie Yan, Sifei Liu, Zhen Lei, Dong Yi, and Stan Z Li. A face antispoofing database with diverse attacks. In *2012 5th IAPR international conference on Biometrics (ICB)*, pages 26–31. IEEE, 2012. [6](#)

- [75] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Xuequan Lu, Ran Yi, Shouhong Ding, and Lizhuang Ma. Instance-aware domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20453–20463, 2023. [1](#), [2](#), [6](#), [7](#), [8](#)
- [76] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Ran Yi, Shouhong Ding, and Lizhuang Ma. Adaptive mixture of experts learning for generalizable face anti-spoofing. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6009–6018, 2022. [6](#), [7](#)
- [77] Qianyu Zhou, Ke-Yue Zhang, Taiping Yao, Ran Yi, Kekai Sheng, Shouhong Ding, and Lizhuang Ma. Generative domain adaptation for face anti-spoofing. In *European Conference on Computer Vision*, pages 335–356. Springer, 2022. [2](#)
- [78] Yixuan Zhou, Yi Qu, Xing Xu, and Hengtao Shen. Imbsam: A closer look at sharpness-aware minimization in class-imbalanced recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11345–11355, 2023. [4](#)
- [79] Zhipeng Zhou, Lanqing Li, Peilin Zhao, Pheng-Ann Heng, and Wei Gong. Class-conditional sharpness-aware minimization for deep long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3499–3509, 2023. [4](#)
- [80] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha C Dvornek, James s Duncan, Ting Liu, et al. Surrogate gap minimization improves sharpness-aware training. In *International Conference on Learning Representations*, 2021. [2](#), [3](#), [4](#), [7](#)