

Open-Set Face Recognition with Maximal Entropy and Objectosphere Loss

Rafael Henrique Vareto^{a,b}, Yu Linghu^b, Terrance Edward Boulton^c, William Robson Schwartz^a, Manuel Günther^{b,c}

^aSmart Sense Laboratory, Federal University of Minas Gerais, Belo Horizonte, Brazil – {rafaelvareto,william}@dcc.ufmg.br

^bDepartment of Informatics, University of Zürich, Zürich, Switzerland – {linghu,guenther}@ifi.uzh.ch

^cVision and Security Technology Lab, University of Colorado Colorado Springs, Colorado Springs, USA – tboulton@vast.uccs.edu

Abstract

Open-set face recognition characterizes a scenario where unknown individuals, unseen during the training and enrollment stages, appear on operation time. This work concentrates on watchlists, an open-set task that is expected to operate at a low False Positive Identification Rate and generally includes only a few enrollment samples per identity. We introduce a compact adapter network that benefits from additional *negative* face images when combined with distinct cost functions, such as Objectosphere Loss (OS) and the proposed Maximal Entropy Loss (MEL). MEL modifies the traditional Cross-Entropy loss in favor of increasing the entropy for negative samples and attaches a penalty to known target classes in pursuance of gallery specialization. The proposed approach adopts pre-trained deep neural networks (DNNs) for face recognition as feature extractors. Then, the adapter network takes deep feature representations and acts as a substitute for the output layer of the pre-trained DNN in exchange for an agile domain adaptation. Promising results have been achieved following open-set protocols for three different datasets: LFW, IJB-C, and UCCS as well as state-of-the-art performance when supplementary negative data is properly selected to fine-tune the adapter network.

Keywords: Neural networks, biometrics, classification, face recognition, open-set, watchlist.

1. Introduction

In open-set face recognition, there is no guarantee that a person caught on camera has been previously enrolled in the gallery of known individuals. Within the open-set task, there are *watchlists*, a scenario that must operate at a very low False Positive Identification Rate (FPIR) as a foresight that the majority of queried individuals are not expected to be registered in the gallery set. When a detected face is mistakenly assigned to one of the identities, it raises a false alarm (a false positive identification) that usually triggers human actions and, therefore, must be avoided to decrease both operational cost and personal discomfort of innocent citizens [1, 2]. Additionally, subjects of interest may be either missed by a face detector or erroneously classified as unknown individuals or assigned a different identity.

Face biometric systems using deep convolutional neural networks have matured into an age of ubiquitous deployment and high performance in recent years. However, most researchers have left open-set problems aside and channeled their efforts into closed-set identification and verification applications. Recently, an outstanding vendor of face recognition technology suffered considerable criticism for matching USA congress members to mugshots of criminals [3]. The incident became an eye-opener on the risks of such commercial identification systems as false alarms can substantially bias security personnel while increasing the responsibility for officers to thoroughly verify the results of the surveillance system. After all, no one would

be contented with innocent people being held up by law-enforcement agencies due to a biometric system error.

Neural networks are biased toward the data they have been trained on and rarely work well with unknown classes. Fig. 1, adapted from Dhamija *et al.* [5], illustrates such behavior on a handwritten digit and character recognition task. Charts (a) and (b) demonstrate that unknown samples (gray dots) cover most of the known classes when the cross-entropy loss is employed, which proved to be insufficient for open-set problems. Contrarily, adopting a cost function that duly handles negative samples attains better class separation and achieves superior performance.

Although the illustration 1(c) may hold true for elementary problems holding abundant samples and very few classes, it is not guaranteed that such behavior would propagate to more demanding biometric applications [2, 6]. In favor of investigating neglected real-world face problems, this study evaluates how open-set loss functions assist neural networks when the training data consists of a few instances per identity. We propose Maximal Entropy Loss (MEL)¹, a function that adds a penalty margin to known identities and increases the entropy for negative samples as it guides a network into differentiating unknown from known subjects. We also implement with an adapter network that is quickly trained by inputting deep features obtained with leading face architectures and avoid retraining deep backbones every time the gallery set is updated.

¹OpenLoss Package: <https://pypi.org/project/openloss/>

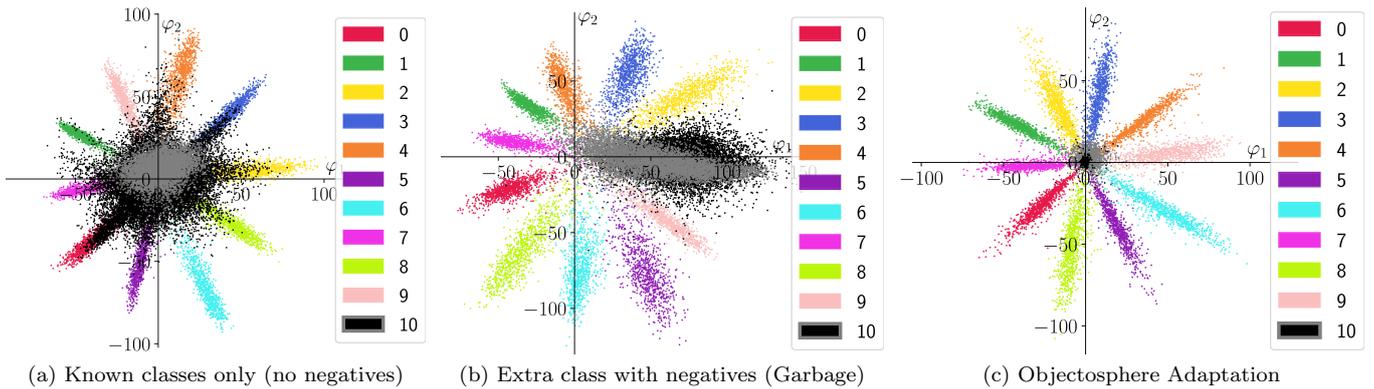


Figure 1: BOOSTING UNKNOWN DETECTION WITH NEGATIVE SAMPLES. The behavior of three different approaches when trained with additional negative data and evaluated with unknown samples. LeNet++ network [4] topologies are trained on 10 MNIST classes (knowns, colored dots) and evaluated with EMNIST letters (negatives, black) as well as Devanagari letters (unknowns, gray).

This work discloses how a compact adaptation network, equipped with few fully-connected layers, responds to open-set protocols on three different datasets, namely LFW, IJB-C and UCCS [7–9]. We exploit data that do not require domain adaptation to perform gallery specialization, in which the knowledge obtained from networks pre-trained on large face datasets is reused to boost performance on related face recognition tasks. We evaluate three architectures for feature extraction, including AFFFE [10], a deep-feature extractor adapted to handle misaligned and blurry faces; VGGFace2 SEnet50 [11], a backbone that takes advantage of its squeeze-and-excitation blocks; and ArcFace [12], a ResNet-101 network that applies a special loss for producing better-suited face representations.

The proposed approach actually differs from most investigations available in the literature. To the best of our knowledge, no genuine open-set face recognition work has been evaluated on the IJB-C benchmark. Most methods typically aim at improving open-set recognition by providing better feature embeddings for face verification, which comprises a different biometric task. Moreover, MEL is the first loss function to simultaneously penalize known and negative samples. Its distinctiveness drives the network toward learning more discriminative face embeddings as it meticulously searches for enhanced parameters.

The major contributions of our work are:

- (a) We evaluate distinct cost functions as well as propose MEL, a novel loss function that maximizes the entropy in order to make training more rigorous.
- (b) We further analyze the Objectosphere loss [2] in favor of verifying how it modifies the feature vector norm of training and test face samples.
- (c) We present an adapter network that accelerates the computationally-expensive retraining or fine-tuning of deep convolutional neural networks.
- (d) We conduct a detailed open-set analysis of all evaluated cost functions on datasets containing thousands of identities but few samples per class.
- (e) We run experiments to verify whether the proposed

approach is effective when combined with distinct deep feature extractors and evaluated on well-known open-set face recognition datasets.

The remainder of this work is organized as follows: Section 2 provides an overview of related work. Section 3 describes the proposed approach: a compact adaptation network combined with MEL or other open-set loss functions. Section 4 exposes the experimental evaluation on three different face datasets. Section 5 presents a thorough discussion of the attained results and Section 6 finishes up with the conclusion and final words.

2. Related Work

Most modern face recognition systems rely on deep convolutional neural networks (DNNs) [12–19]. Strategies have been designed to achieve better identification performance on difficult images, such as margin-based or triplet loss, and different network topologies [11, 20, 21]. However, DNNs are not usually designed to handle facial images with a low optical resolution, or even false-positive face detections. Besides, the aforementioned works do not “disregard” low-interest samples and, as a result, end up matching all unknown identities with their respective most similar subjects from the gallery set.

Vareto *et al.* [22] combined hashing functions to set up a vote-list histogram. Some researchers have adopted *one-vs-all* SVM or PLS models [23, 24] whereas others explored clustering techniques [25, 26]. The aforementioned methods neither implement the entire closed-set identification pipeline nor comply with the requirements of real-time or real-world applications. Most present-day methods aim at improving closed-set recognition or person re-identification problems and rarely consider open-set protocols [15–17]. Others typically focus on open-set recognition by providing better feature embeddings for face verification, which comprises a different biometric task [12, 18, 19]. Moreover, Hassen *et al.* [27] introduced a loss function that draws same-class samples near and Zhou *et al.* [28] introduced

an additional layer to store class-specific thresholds. Researchers have also explored adversarially-generated samples for “balanced” decision boundaries among known and unknown classes [29–31]. However, these approaches have been evaluated on datasets holding numerous samples per class and, as a consequence, they are not an accurate portrayal of real-world biometric problems.

Most used datasets in non-face open-set recognition are CIFAR [32], MNIST [33], SVHN [34] and TinyImageNet, a subset of ImageNet [35], to name a few. They range from 5 to 20 classes in the known set, but each class encompasses myriads of samples. Approaches evaluated on such data are not hampered by the shortage of image samples available for training and, in fact, better preserve the inherent data distribution [36]. Labeled Faces in the Wild (LFW) [7] used to be the leading facial benchmark. LFW contains 13,233 images unevenly distributed among almost six thousand classes. As it was initially designed for verification, experts have proposed non-official open-set protocols [37, 38]. IJB-C [8] contains two disjoint gallery partitions of known individuals merged together for closed-set recognition. The open-set protocol requires the use of a single gallery partition and, hence, half of the probe subjects have no corresponding match in the gallery set.

In contrast to IARPA’s benchmarks, the original Uncontrolled College Students (UCCS) dataset [39] and its extended version [9] mandate for faces to be detected as part of the recognition pipeline. The UCCS dataset consists of images captured at a university campus covering different weather conditions. UCCS’s gallery set encompasses 1,085 known subjects, with approximately 20 instances per class, and countless face samples not labeled to any of the known identities. There are several partially-occluded faces due to lamp posts and tree branches along with accessories like sunglasses, hats, hoodies, or fur jackets that make both detection and recognition in UCCS benchmark a challenging task.

In summary, few works have designed methods to properly tackle open-set face recognition with mechanisms that enable the network to differentiate individuals of interest from unknown people in a scenario with thousands of identities but few samples per class. With that in mind, we evaluate our proposed approach on realistic face datasets as a meaningful contribution to the biometric discipline. Due to these fundamental properties and their intrinsic open-set nature, we use both IJB-C and UCCS datasets along with LFW in our experiments.

3. Proposed Approach

A watchlist application S generally consists of three sequential stages: $S = S_d \rightarrow S_r \rightarrow S_c$ and should raise an alarm only when probe samples belong to gallery set G . Subsystem S_d corresponds to the face detection and landmark localization method locates faces in the original input image. For every detected face, the representation

module S_r extracts a corresponding numerical feature vector. The identification subsystem S_c assigns one of the gallery identities $g \in G$ to the probe face sample. As shown in Fig. 2, we introduce an additional adaptation module S_a that takes original features from the representation stage and further transforms them into attributes that are better suited for the task at hand.

Template $T_g = S_a(S_r(S_d(x_g)))$ corresponds to the mean representation of subject g when multiples sample are available per class. Similarly, $F_p = S_a(S_r(S_d(x_p)))$ becomes the probe representation. The classification subsystem S_c computes a similarity score $s(T_g, F_p)$ between F_p and template T_g for each known individual $g \in G$. Then, S_c rejects probe samples as unknown when they attain scores lower than θ for every subject of interest. If not, F_p is assigned to the identity holding highest score $\max_{g \in G} s(T_g, F_p)$.

The compact adapter network S_a aims to establish a drastic difference between gallery subjects and unknown faces. Therefore, it is not possible to enroll new subjects in the gallery set without retraining. Since we rely on features extracted from representational network S_r , retraining the adapter network S_a is fast and can be performed whenever a new subject needs to be enrolled – given that watchlists are oftentimes relatively stable over time.

3.1. Training

One of the fundamental aspects behind the procedure depicted on Fig. 2 is that any pre-trained network, such as VGGFace2 [11], AFFFE [10] and ArcFace [12] can be adopted as the pipeline’s face representation subsystem S_r . Consequently, the proposed approach does not require time-consuming retraining of massive deep networks every time a new subject is inserted into gallery G since a small adapter network S_a fits the extracted set of representations $R_g = S_r(x_g) \forall g \in G$.

Adapter Network. The adapter network S_a consists of a multi-layer perceptron network with fully-connected layers. In fact, S_a is composed of an input layer L^i , two hidden layers L^{h1} and L^{h2} , and an output layer L^o . The input layer takes in feature vectors R extracted with a pre-trained DNN S_r and, therefore, its size varies according to the deep feature dimension. The first hidden layer L^{h1} incorporates a non-linear hyperbolic tangent activation function that outputs values in the range -1.0 to $+1.0$ whereas L^{h2} delivers a compact feature representation.

The learning strategy is similar to the training process followed by traditional face recognition systems: we set the output layer L^o to hold a size analogous to the number of gallery-enrolled identities. In other words, each last layer’s logit node $L_g^o \in L^o$, also denoted as $l_g \in L^o$, stands for the corresponding activation of known subject $g \in G$. In general, these activations are employed for open-set face classification but they present inferior performance when compared to the distance computation of deep features obtained with neural networks [2].

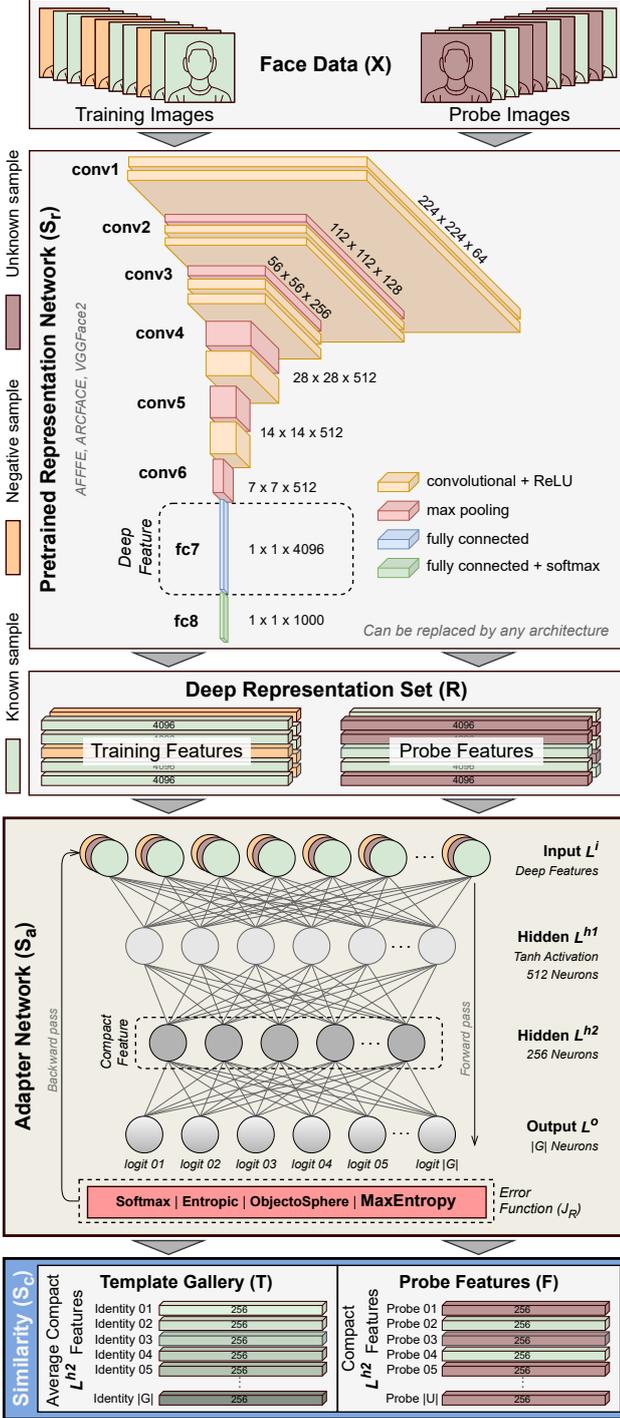


Figure 2: PROPOSED PIPELINE. Given any set of images, training features comprising known and negative samples are input to the proposed adapter network S_a for learning parameters that minimize the adopted loss function. Compact 256-dimensional features from the penultimate layer of S_a are averaged to build a gallery of templates (T) during enrollment. On testing, compact features are extracted from probe data and compared with gallery templates T through cosine similarity.

The adapter network has been originally designed as a conventional multi-layer perceptron network. Ordinarily, its output logit layer L^o could be associated with the

Softmax activation function A_S and assumes a role as the ultimate recognition phase:

$$A_S(l_g) = \frac{e^{l_g}}{\sum_{1 \leq g' \leq |G|} e^{l_{g'}}} \quad (1)$$

However, it performs differently considering that S_a also provides discriminative feature representations that are required in the subsequent similarity classification subsystem S_c (see the blue rectangle in Fig. 2). The adapter network yields its two last layers during the training stage: logits from L^{h2} input ObjectSphere while L^o values feed the remaining loss functions as detailed below.

Entropic Open-Set Loss (J_E) [5]. The Entropic Open-set loss comes to maximize the uncertainty of negative samples by inducing their Softmax responses to lie uniformly distributed. J_E boosts the maximum entropy distribution of uniform probabilities from negative samples over all $|G|$ known classes registered in the gallery set G . In the classic Cross-Entropy loss, t_g represents a one-hot vector holding the value of one at the index that corresponds to known class g . Under the inclusion of negative instances, J_E attributes uniform values to target vector $\forall g : t_g = \frac{1}{|G|}$ in such a way that unseen samples are considered as equal members of each known identity:

$$J_E = - \sum_{1 \leq g \leq |G|} t_g \log A_S(L_g^o(R_x)) \quad (2)$$

Maximal Entropy Loss (J_M). The proposed Maximal Entropy loss associates the previously stated Entropic Open-set loss with margin-based Softmax (A_{S_m}) [40, 41]. Equation (3) points out how A_{S_m} affixes a non-negative penalty margin m to A_S in order to decrease the intra-class distance and maximize the segregation among distinct classes. As the penalty increases, a network learns parameters that push samples more firmly toward their class centroids. The parameter defines a distance among different classes and, consequently, draws same-class samples closer [40].

$$A_{S_m}(l_g) = \frac{e^{l_g - m}}{e^{l_g - m} + \sum_{g' \neq g} e^{l_{g'}}} \quad (3)$$

The Maximal Entropy Loss J_M combines the best of both worlds since the Soft-Margin Softmax targets known training samples whereas the Entropic Open-set handles negative instances available during the learning stage. More precisely, function J_M maximizes the entropy regarding the correct target class when $x \in G$ in the interest of making the closed-set identification more rigorous and, as a result, equips the adapter network with more discriminative weights. The handicap parameter m establishes a decision boundary for a more appropriate separation of known individuals:

$$J_M = \begin{cases} -\log A_{S_m}(L_g^o(R_x)) & \text{if } x \in G \\ -\frac{1}{|G|} \sum_{g=1}^{|G|} \log A_S(L_g^o(R_x)) & \text{if } x \notin G \end{cases} \quad (4)$$

For a negative sample $x \notin G$, the designed loss uniformly distributes the target variable score among all $g \in G$ subjects in an attempt to support the network in distinguishing gallery-enrolled subjects from unknown identities. Similar to the aforementioned J_E loss, the insight of equalizing logit values for unknown samples lies behind not knowing anything about their corresponding identity and, therefore, they hold an equivalent likelihood of being assigned to any subject registered in the gallery set. Analogous to Dhamija *et al.* [5], the overall error obtained with J_M is minimized when the Softmax responses $A_S(\cdot)$ of negative samples are equally distributed.

Objectosphere Loss (J_O) [5]. Objectosphere dissociates representations of known and negative samples by directly modifying their feature magnitudes. Since J_E cannot guarantee that such a pattern would be generated for unknown samples, Fig. 1(c) illustrates that Objectosphere modifies the network weights to drive negative instances toward the feature space origin. This is achieved by forcing the magnitude of negative features $\|L^{h2}(R_x)\|_2$ to be closer to zero while simultaneously pushing known feature magnitudes to at least ξ , a required hyperparameter for Objectosphere.

$$J_O = J_E + \lambda \begin{cases} \max(\xi - \|L^{h2}(R_x)\|_2, 0)^2 & \text{if } x \in G \\ \|L^{h2}(R_x)\|_2^2 & \text{if } x \notin G \end{cases} \quad (5)$$

Larger ξ values scale up deep features, including those extracted from unknown samples, which can be compensated by lower weights in the last layer L^o ; however, what actually makes a difference is the increased separation among known, negative and, ultimately, unknown samples.

Additional Garbage Class (J_G). With the high demand for open-set recognition systems and the practicability of the Cross-Entropy loss, accessible in every deep learning framework, a common strategy is to add an extra class $|G| + 1$ to encompass negative samples. We refer to the adapter network S_a trained with J_G as the **Garbage** approach in the experimentation section.

$$J_G = - \sum_{1 \leq g \leq |G|+1} t_g \log A_S(L_g^o(R_x)) \quad (6)$$

3.2. Enrollment and Inference

The enrollment of subjects of interest is illustrated in Fig. 2. It starts with the extraction of compact features from all gallery samples in the interest of creating a gallery of templates T . Equation (7) demonstrates that for each known identity $g \in G$, a unique template T_g is established by averaging the normalized compact features obtained with the adapter network where $|K_g|$ is the number of enrollment samples available for subject g .

$$T_g = \frac{1}{|K_g|} \sum_{1 \leq k \leq K_g} S_a(S_r(S_d(x_{g,k}))) \quad (7)$$

Analogous feature vectors are obtained for probe images $x_p \in P$ during the inference stage by employing the very same representational and adaptation networks utilized in the enrollment phase:

$$F_p = S_a(S_r(S_d(x_p))) \quad (8)$$

Then, the classification module S_c computes similarity scores between probes and all gallery-enrolled identities through the angular cosine similarity:

$$s(T_g, F_p) = \cos(T_g, F_p) = \frac{T_g^T F_p}{\|T_g\| \cdot \|F_p\|} \quad (9)$$

It is worth mentioning that we have also investigated other similarity-based functions that make use of probe feature magnitudes [2, 5]; however, they include several issues that have not been addressed in this work.

4. Experiments

This section presents the experimental evaluation of the approaches described in Section 3. It starts detailing the adopted evaluation metrics, assessed methods, and a description of the experimental setup along with the explored datasets. Further, it provides an experimental assessment of the obtained feature magnitudes and a comparison between the traditional Cross-Entropy and the negative-based cost functions, namely Entropic Open-set, Objectosphere, and the proposed Maximal Entropy Loss.

4.1. Evaluation Metrics

We adopt the open-set ROC curve [42–44], which plots the True Positive Identification Rate (TPIR) against the False Positive Identification Rate (FPIR) by varying the rejection threshold θ . TPIR is computed solely on probe samples of known subjects \mathbf{K} by considering probes to be correctly identified if the similarity to the correct identity g^* is the highest and above operating threshold θ :

$$\text{TPIR}(\theta) = \frac{1}{|\mathbf{K}|} \left| \left\{ F_p \in \mathbf{K} \mid \arg \max_{g \in G} \cos(T_g, F_p) = g^* \wedge \cos(T_{g^*}, F_p) \geq \theta \right\} \right| \quad (10)$$

FPIR corresponds to the false alarm rate triggered by unknown samples \mathbf{U} . A false positive identification occurs when the similarity of an unknown sample F_p to any of the known subject templates T_g is larger than threshold θ :

$$\text{FPIR}(\theta) = \frac{1}{|\mathbf{U}|} \left| \left\{ F_p \in \mathbf{U} \mid \max_{g \in G} \cos(T_g, F_p) \geq \theta \right\} \right| \quad (11)$$

An optimal open-set face identification system presents a TPIR of 1 at an FPIR of 0. By varying the threshold θ , the open-set ROC curve can be created.

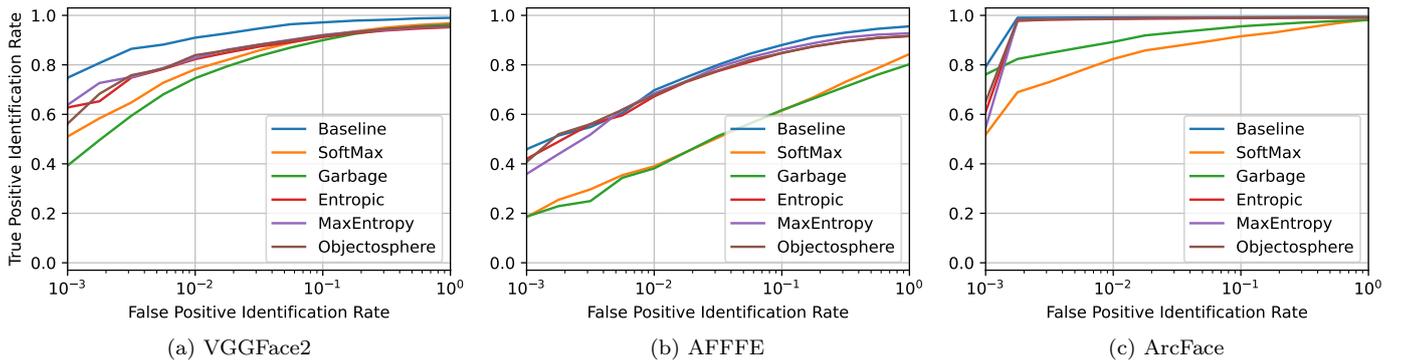


Figure 3: LFW EVALUATION. Open-set ROC charts are shown for VGGFace2, AFFFE, and ArcFace features. Due to the small size of the LFW dataset, FPIR values smaller than 10^{-3} cannot reliably be computed and are, hence, left out. Because of the small amount of three training samples per identity, the adapter network is not able to provide more meaningful features than the representation network S_r .

4.2. Evaluated Datasets

We utilize a data partition [26, 37] that splits LFW into three disjoint groups: 602 known, 1070 negative, and 4096 unknown identities. We use the provided hand-labeled landmarks available in LFW dataset during the alignment process. For IJB-C, we train the method on *gallery A* only so that all *gallery B* matching identities available in the probe set act as unknown face samples. Additionally, LFW is incorporated as the negative set since none of its classes are encountered in IJB-C. UCCS metadata provides bounding boxes and identity labels, containing either known subject identities or negative labels for unknown faces. We incorporate the MTCNN face detector [45] as the default detection system S_d on IJB-C and UCCS benchmarks. We employ the very same face detector throughout the experiments to standardize the face detection stage. Following the evaluation protocol, all background detections of MTCNN serve as additional unknown samples during testing in the UCCS dataset [9].

4.3. Evaluated Approaches

In the interest of comparing the proposed adapter network along with Maximal Entropy and Objectosphere loss functions to other methods, we incorporate four additional approaches: **Baseline**, **SoftMax**, **Garbage** and **Entropic**. Apart from **Baseline**, all evaluated methods run the complete pipeline depicted in Fig. 2 in which the template gallery consists of feature vectors extracted from the adapter network S_a . In addition to the adapter network, we also investigate whether it is beneficial to fine-tune the entire feature backbone model on the gallery data, which has been shown to be beneficial for larger datasets. Since this training is much more time consuming, we restrict our experiments to the largest and most difficult dataset, *i. e.*, IJB-C. The seven evaluated techniques are:

- **Baseline** consists of creating a template set with the original features extracted from the representational system S_r and computing the cosine similarity.

- **SoftMax** follows the proposed pipeline by training the adapter network S_a with Cross-Entropy loss, without exploiting any negative samples (negative-free).
- **Garbage** extends **SoftMax** as it creates a template T_g for each known individual $g \in G$ along with an exclusive template $T_{|G|+1}$ holding negative samples.
- **Entropic** also follows the proposed pipeline, but this time the Entropic Open-set loss is adopted to handle known and negative samples (negative-based).
- **Objectosphere** adopts the Objectosphere loss to train the adapter network with hyperparameters $\xi = 1$ and $\lambda = 0.01$, as specified in Equation (5).
- **MaxEntropy** consists of training the adapter network S_a with the proposed Maximal Entropy loss, which holds hyperparameter $m = 0.40$ as a default value.
- **Finetuning** involves training all layers of the adopted architecture on the evaluated IJB-C dataset.

4.4. Network Setup

The network S_a benefits from representation systems S_r , that is, AFFFE, ArcFace and VGGFace2 [10–12] with 1000, 512 and 2048-dimensional deep features, respectively. The feature extraction counts on Bob’s [46, 47] biometric pipeline² that handles face detection, alignment and feature extraction. The adapter network S_a topology is a compact fully-connected network with 512 and 256 neurons in the two hidden layers. Given the aforesated hyperparameters, VGGFace2 composes the worst-case scenario in which the adapter network would hold no more than 1.7 million trainable weights, which corresponds to a small fraction of the total of 138 million parameters contained in the deep backbone (98% less than VGG-16).

The pipeline is built upon *PyTorch* framework and consists of 500 training epochs for all datasets. Convergence

²<https://www.idiap.ch/software/bob/docs/bob/docs/stable/>

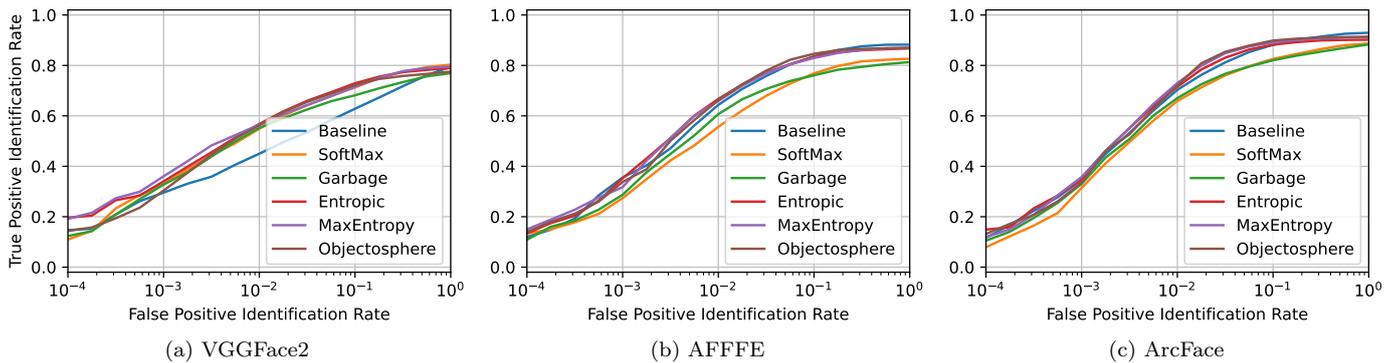


Figure 4: UCCE EVALUATION. *UCCE is a more challenging dataset than LFW as it comprehends a surveillance and unrestricted domain. Therefore, training the adapter network using UCCE known and negative samples improves the performance over the Baseline, especially when training and evaluation samples hold equivalent distribution.*

on the validation set was commonly achieved in the first 150 epochs, only minor improvements have been encountered after 200 epochs. When disregarding the feature extraction process performed in S_r , the training procedure takes around 20 minutes for LFW, 80 minutes for UCCE and no more than three hours for IJB-C on a regular multi-core desktop computer with a single NVIDIA Titan X GPU. If more training speed is required, the network topology can be adapted, the number of epochs can be reduced or more GPU resources can be added.

4.5. Comparison to the State of the Art

In the interest of showing the advantage of **MaxEntropy** and **Objectosphere** over **SoftMax**, **Garbage** and **Entropic**, the adapter network S_a is trained on different face datasets holding the very same topology and hyper-parameters for all dependent methods. Figures 3 through 6 depict several approaches in which all of them, except **Baseline**, rely on S_a . Additionally, Tab. 1 provides a detailed list of TPIR values for selected FPIR operating points, evaluated on all three network topologies and all three datasets. The results obtained on the three evaluated datasets are described in the following paragraphs:

Labeled Faces in the Wild. Fig. 3 portrays the investigation on LFW considering different feature representations: VGGFace2, AFFFE and ArcFace. **Baseline** presents an outstanding performance using VGGFace2 representation module in Fig. 3(a), implying that no supplementary data is required for LFW due to its innate characteristics. Plots (b) and (c) point out a comparable performance between **Baseline** and negative-based cost functions.

There is an equivalent behavior with AFFFE when the false-positive proportion exceeds three per thousand samples (3×10^{-3}). ArcFace backbone equipped all approaches with discriminative feature vectors so that very little can be concluded in terms of accuracy. Note that four methods attained open-set performance greater than 95% in (c) when FPIR surpasses 2×10^{-3} . However, results are substantially inferior under **SoftMax** or **Garbage** approach.

Unlike most recent face datasets, LFW consists of reasonably good-quality images of cropped faces that cooperate with deep networks in delivering satisfactory feature representations. As a consequence, computing the cosine distance among original feature vectors, as performed by **Baseline**, is sufficient to go toward the state of the art. The small amount of data (three images per subject) seems insufficient to train the adapter network with traditional cost functions. The adopted non-official protocol [37] holds nearly 9,300 samples in the probe set and, therefore, the actual threshold value is estimated at no more than 10 images when the FPIR is less than 10^{-3} . Moreover, the TPIR performance score for scarce samples is not reliable in low FPIR regions due to the natural threshold fluctuation.

UnConstrained College Students. Fig. 4 discloses the experimental evaluation on the UCCE benchmark. Along with identities composing the gallery set, UCCE data encompasses both false positive detections (misdetected) and faces from unknown subjects. **MaxEntropy** seems capable of attenuating the domain difference between the source data used to train the representation network S_r and the student population scope present in the UCCE dataset. On the other hand, the domain adaptation seems less impactful for ArcFace features, which indicates that ArcFace architecture can be used in various domains.

Fig. 4(a) reveals that our approach can benefit from the addition of negative samples as the best overall result was achieved with the adapter network when trained with the proposed Maximal Entropy loss. Fig. 4(b) also signals significant accuracy gained through the addition of negative samples. The chart indicates that AFFFE face representations are better adapted for low-resolution images than VGGFace2; however, both are surpassed by ArcFace’s robust feature vectors. Fig. 4(c) shows that the negative-exploring cost functions obtain analogous performance: slight dominance of **MaxEntropy** when FPIR is between 10^{-3} and 10^{-1} . Although **Baseline** prevails in the interval $[10^{-1}, 10^0]$, it attains lower accuracy in the aforementioned range along with the other methods.

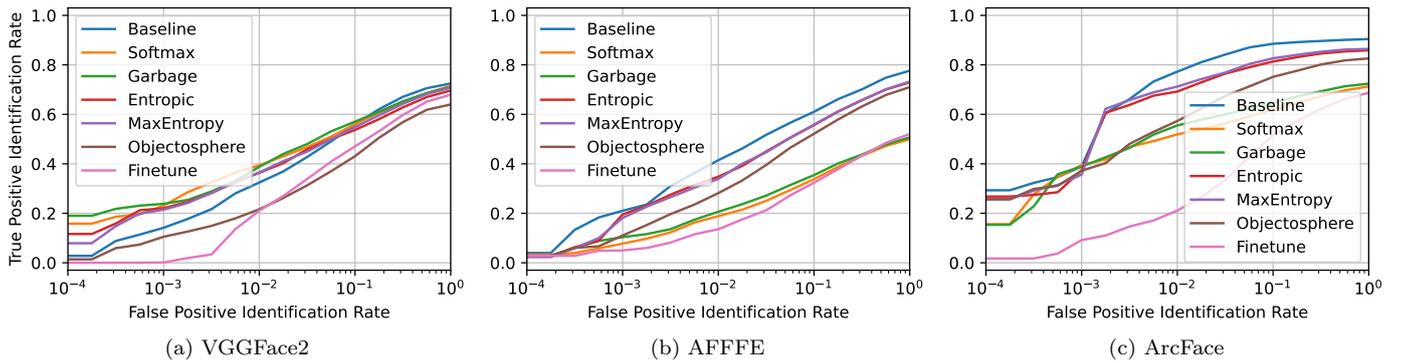


Figure 5: IJB-C + LFW EVALUATION. Open-set ROC charts are shown for AFFFE, ArcFace and VGGFace2 features. This evaluation follows IJB-C’s open-set protocol test 4 with the addition of the entire LFW dataset as negative samples. Negative data diverging from the gallery set distribution seem incapable to contribute to the method’s performance.

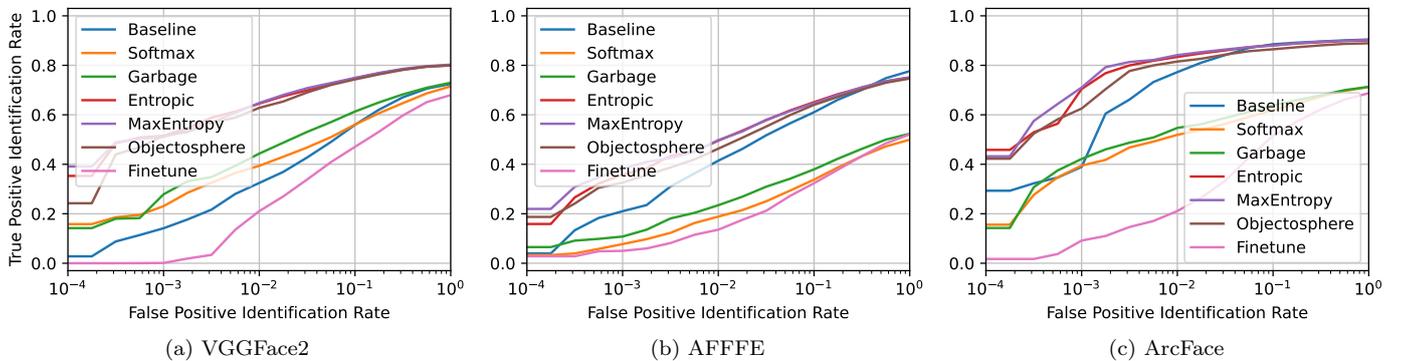


Figure 6: IJB-C EVALUATION. Open-set ROC charts are shown for AFFFE, ArcFace and VGGFace2 features. Negative samples are obtained from gallery set B of IJB-C dataset. This evaluation does not adhere to IJB-C’s open-set protocol test 4. When negative samples embody the same distribution as known samples, “negative-based” cost functions along with the adapter network outperform the Baseline.

IARPA Janus Benchmark C series. Fig. 5 exposes experiments on IJB-C merged with more than 13,000 negative samples acquired from LFW. The discrepancy³ in image resolution and pose variations between both datasets ends up reflecting on the results as LFW does not play a decisive enhancement role in the proposed adapter network’s identification performance when assessing IJB-C benchmark. The three plots suggest that the Finetune approach could not maintain the generalization capability of the original backbone performance either combined with cosine similarity (Baseline) or the adapter network.

According to Fig. 5(a), negative samples do not seem to provide significant improvement when evaluating VGGFace2 feature vectors and, in fact, they turn out to impair Objectosphere’s exactness. Fig. 5(b) corresponds to experiments containing AFFFE features and shows that Baseline outperforms all other approaches. MaxEntropy attains comparable performance at a low false positive identification rate when it ranges from 1×10^{-3} to 3×10^{-3} .

³IJB-C contains images without standardized traits whereas LFW comprises mostly good-quality images with close-to-frontal faces.

ArcFace experiments in Fig. 5(c) also demonstrate the dominance achieved with the Baseline approach. An approximate accuracy is reached by the MaxEntropy method when FPIR comprises the area to the left of 2×10^{-3} .

5. Discussion

This section examines the effect of training the adapter network with different-distribution data: IJB-C consists of lower-resolution images with a wider range of poses and expressions when compared to LFW. Tab. 1 provides a complete view of the results for different FPIRs.

5.1. Differences between IJB-C and LFW

Dhamija *et al.* [5] pointed out that the choice of negative samples plays an important role when training an open-set network. Fig. 8(a) shows that LFW does not follow the same feature distribution as IJB-C. As revealed in Fig. 5, selecting LFW to compose the set of negative samples could not provide further improvements and outperform the baseline method on IJB-C dataset except for experiments containing VGGFace2 representations.

Table 1: OPEN-SET ROC EVALUATION. *Open-set ROC results are shown for AFFFE, ArcFace and VGGFace2 feature representations on the three evaluated datasets, namely LFW, UCCS and IJB-C. Each cell consists of the True-Positive Identification Rate where False-Positive Identification Rate value is indicated in the first column (TPIR@FPIR). Best values per model are highlighted in bold, second in italics.*

O-ROC	Datasets:		LFW			UCCS			IJB-C		
	FPIR	Methods	VGG2	AFFFE	ArcFace	VGG2	AFFFE	ArcFace	VGG2	AFFFE	ArcFace
10^0	Baseline	0.99128	0.99128	0.99631	0.79398	0.88193	0.92892	0.72426	0.77669	0.90307	
	SoftMax	<i>0.96887</i>	<i>0.96887</i>	0.98649	0.80361	0.82651	0.88795	0.68137	0.47853	0.66994	
	Garbage	0.96264	0.96264	0.98034	0.76747	0.81325	0.88313	0.67525	0.40368	0.61595	
	Entropic	0.95268	0.95268	0.99631	0.79036	0.86506	0.90000	0.69240	<i>0.73374</i>	0.85521	
	Objectosphere	0.95143	0.95143	0.99631	0.77831	<i>0.87108</i>	0.91205	0.67034	0.70920	0.82699	
	MaxEntropy	0.95641	0.95641	0.99631	<i>0.79639</i>	<i>0.87108</i>	<i>0.91566</i>	<i>0.70711</i>	0.73006	<i>0.86626</i>	
10^{-1}	Baseline	0.97136	0.87975	0.99386	0.62651	0.83253	0.88313	0.56127	0.61104	0.88466	
	SoftMax	0.91158	0.61595	0.91646	0.71084	0.76988	0.82771	0.52574	0.31043	0.57301	
	Garbage	0.89788	0.61595	0.95577	0.68193	0.76145	0.82048	0.52083	0.23926	0.49939	
	Entropic	0.91283	0.84540	0.99017	0.73253	0.83012	0.87952	0.53064	0.55092	0.81227	
	Objectosphere	<i>0.92030</i>	0.84540	0.99140	0.71807	0.84819	0.90000	0.47181	0.51656	0.73865	
	MaxEntropy	<i>0.92030</i>	<i>0.86258</i>	0.99386	<i>0.71928</i>	<i>0.83976</i>	<i>0.89398</i>	<i>0.54289</i>	<i>0.55951</i>	<i>0.82454</i>	
10^{-2}	Baseline	0.90909	0.69693	0.99509	0.44940	0.64337	0.70120	0.32230	0.41227	0.77178	
	SoftMax	0.78207	0.39018	0.82310	0.54699	0.55783	0.66024	0.32598	0.17546	0.44908	
	Garbage	0.74595	0.38160	0.89312	0.55060	0.60482	0.67108	0.33211	0.13129	0.38773	
	Entropic	0.82067	0.66748	0.98526	0.57229	0.66988	0.71325	0.37745	0.33742	0.69202	
	Objectosphere	<i>0.83935</i>	0.67853	0.98894	0.56386	0.66024	<i>0.72410</i>	0.26348	0.27853	0.53988	
	MaxEntropy	0.82939	<i>0.68466</i>	<i>0.99017</i>	<i>0.56747</i>	<i>0.66627</i>	0.73253	<i>0.37500</i>	<i>0.34479</i>	<i>0.70552</i>	
10^{-3}	Baseline	0.74720	0.45767	0.79238	0.29277	0.35663	<i>0.35060</i>	0.14093	0.21104	<i>0.39387</i>	
	SoftMax	0.51059	0.18405	0.51966	0.33614	0.26988	0.31325	<i>0.22304</i>	0.09816	0.34969	
	Garbage	0.38979	0.18405	<i>0.76290</i>	0.32651	0.28916	0.32651	0.20221	0.05276	0.28221	
	Entropic	0.62889	<i>0.42209</i>	0.61794	<i>0.34337</i>	<i>0.35301</i>	0.34337	0.21814	0.17301	0.36319	
	Objectosphere	0.56663	0.41104	0.66216	0.30120	0.33976	0.33614	0.16299	0.08834	0.38405	
	MaxEntropy	<i>0.64010</i>	0.35583	0.55283	0.35904	0.31687	0.36145	0.23897	<i>0.18650</i>	0.40982	

IJB-C probe samples as well as its enrollment data are distributed differently. More precisely, gallery-enrolled samples contain mostly good-quality still photos whereas probe samples are mainly composed of low-resolution still images or blurred video frames. We tend to believe that the adapter network S_a over-adapts to good-quality enrollment samples when it inputs only high-standard data. Therefore, module S_a ends up lowering the performance on IJB-C by rejecting many probe samples as unknown. Selecting enrollment and probe data with similar distribution is likely to increase performance.

Fig. 6 discloses an additional set of experiments on IJB-C benchmark in which *gallery A* populates the known set and *gallery B* composes the negative set for **Entropic**, **MaxEntropy** and **Objectosphere**. This scenario affords a related data distribution between both training subsets. Despite probe and enrollment data sharing different capture quality, results show that appropriate negative samples significantly improve the open-set face recognition pipeline. All charts indicate a dominance of **MaxEntropy** over negative-free methods when FPIR lies below 10^{-1} . Factually, using ArcFace backbone achieves the highest accuracy rate of all experiments conducted on IJB-C dataset.

We reckon that real-world watchlist applications would scarcely ever contain negative identities overlapping with unknown face samples. However, the assessment displayed in Fig. 6 provides a reference point on the maximal identification correctness. Results show a recurring superiority of **MaxEntropy** regardless of the adopted representation network. Unlike Fig. 5, where gallery and negative samples hold contrasting data distribution, the resemblance

between both IJB-C disjoint galleries delivers discriminative class boundaries. Distribution-alike data is a must-have aspect required by “negative-based” error functions when seeking negative samples in exchange for a meaningful contribution to the open-set recognition pipeline.

Although experiments showed in Fig. 6 do not adhere to the official IJB-C protocol, there are scenarios in which this training scheme would be appropriate. For instance, an enterprise may have premium clients that must be treated differently than regular customers. They could be addressed by name and offered a comfortable room on the premises. Privileged customers constitute known classes but the remaining ones are placed in the negative set. Prospect customers (unknowns) lie somewhere in between and shall be treated better than the ordinary, but not as good as premium. Consequently, the face recognition system is supposed to raise an alert whenever premium customers come over.

5.2. Deep Feature Magnitudes

Objectosphere loss aims to push feature magnitude extracted from unknown samples to very low figures. It simultaneously attempts to shift the magnitude of known samples toward a specified value ξ . Robust open-set methods are expected to achieve high accuracy in different data with consistent parameters. This requirement plays an essential role in biometrics as it is not possible to anticipate the visual traits of all probe samples. Best results have been attained on UCCS when setting **Objectosphere** parameters $\xi = 1$ and $\lambda = 0.01$, and we have verified that these parameters also work well on LFW and IJB-C.

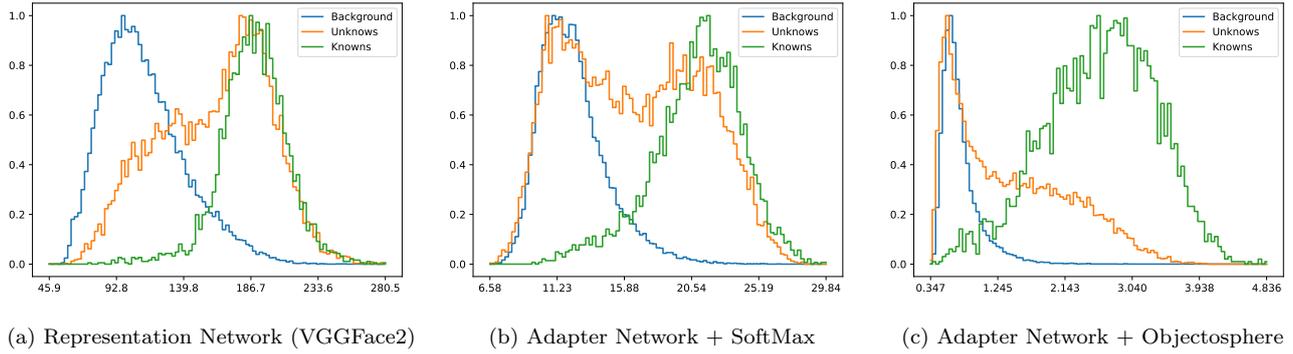


Figure 7: UCCS MAGNITUDES. All plots portray results obtained with probe data: **Knowns** designates subjects registered in the watchlist, **Unknowns** specifies probe samples without corresponding identity in the gallery set, and **Background** refers to face misdetections. Feature magnitude is an indicator that domain adaptation improves separation between known and unknown subjects. While SoftMax provides low magnitudes for background and approximately half of the unknown samples, Objectosphere delivers even better separation.

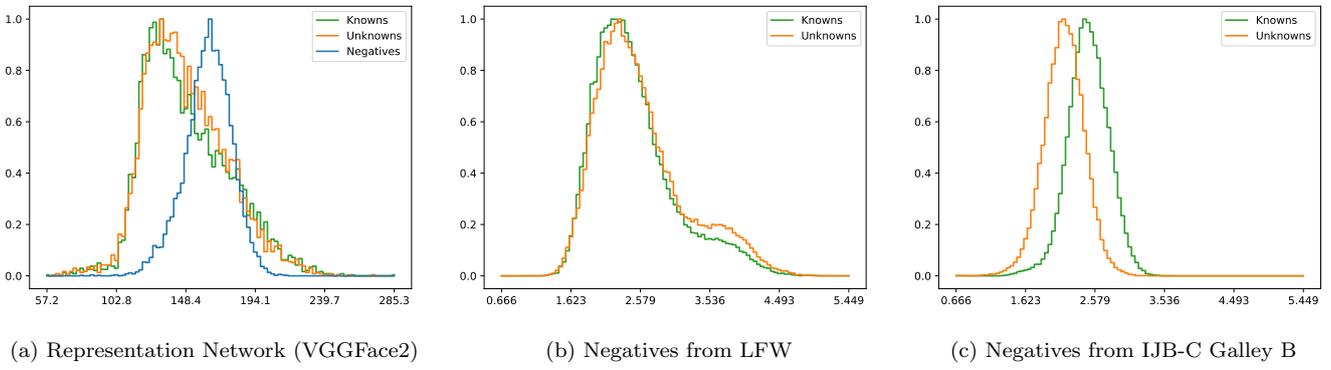


Figure 8: IJB-C AND LFW MAGNITUDES. Chart (a) exposes training feature magnitudes obtained with VGGFace2 in which knowns and unknowns come from IJB-C and negatives derive from LFW. Plots (b) and (c) demonstrate how the adapter network S_a combined with Objectosphere behaves on the evaluation data when trained with negatives coming either from LFW or IJB-C gallery B. Note that gallery B provides better separation between knowns and unknowns whereas LFW is not sufficient to push the distributions apart.

Fig. 7 displays deep feature magnitude histograms for UCCS evaluation data. Original VGGFace2 features hold a considerable magnitude overlap among unknown and known subjects as well as false-positive detections in the background. The intersection remains when training the adapter network with SoftMax, but Objectosphere reduces the coincidental area between known and unknown samples. Basically, weights learned with Objectosphere can distinguish enrolled subjects from unknown identities during the testing stage. Known samples are distributed closer to the desired target magnitude whereas negative (background) samples have a peak close to 0, but are distributed throughout the range of magnitudes.

As indicated in Fig. 8(a), LFW images provide higher magnitudes but IJB-C instances result in low-magnitude representations. Deep networks may misclassify probe samples since the image quality has an impact on the acquired feature vectors. Due to the lack of similarity between IJB-C and LFW, the latter is not capable to guide the adapter network S_a in discriminating IJB-C probe samples. Fig. 8(b) and (c) present probe feature magnitudes when S_a is trained with negatives proceeding from LFW or IJB-C’s gallery B. Note that the magnitudes are well-

above the intended separation threshold $\xi = 1$ and, hence, appropriate negatives might help to separate further.

5.3. Proposed Approach Applicability

MaxEntropy requires a distance margin m in the interval $[0, 1]$ whereas Objectosphere includes sphere-related and regularizing parameters (ξ and λ , respectively). As a result, associating both losses culminates in the specification of three hyper-parameters, not including the ones regarding the adaptation network, such as the number of neurons, learning rate, and batch size to name a few. Cost functions that require the adjustment of multiple parameters make their deployment unfeasible in both academic and realistic scenarios. Consequently, we do not combine MaxEntropy with Objectosphere since a large number of tunable parameters turns into an optimization problem.

We acknowledge that a desirable open-set face recognition approach would only require the enrollment of subjects of interest without the need to fine-tune the deep representation backbone. However, the three evaluated datasets are composed of numerous identities holding very few samples per class in the training set, a common requisite in watchlist problems. Since applying an untouched

pre-trained representation model to dissimilar data distributions regularly results in a substantial accuracy loss, the designed adapter network S_a offers a flexible trade-off between computational time and correctness.

6. Conclusion

Pre-trained deep networks usually require considerable time to be adapted and retrained to new domains, especially when the training data is constantly updated. This is the scenario in which the proposed compact adapter network comes in handy as it serves as a quick-trainable replacement for the output layer. Moreover, the evaluated cost functions take advantage of supplementary information when adding negative samples to the training stage. Experiments have shown that additional samples play an important role in “identifying” the unknown when the training samples are sufficiently representative of the uninvestigated feature space.

The proposed approach is suited for watchlists and transfer-learning tasks since the adaptation network can be attached to the output of any pre-trained deep network model and be quickly adjusted to different data distributions. Retraining large deep backbones, such as ArcFace and VGGFace2, every time a new identity is added to the gallery set becomes categorically infeasible and has proven to be counterproductive. The ArcFace network, for instance, contains nearly 50 million weights in contrast to 394,850 parameters in the adapter network when trained on LFW and inputting 512-dimensional feature embeddings from ArcFace/ResNet-100.

Experiments carried out on the open-set face recognition protocols of LFW, UCCS, and IJB-C have provided a comprehensive analysis of the compact network and the employed loss functions. The evaluation has shown that the association of the adapter network with Objectosphere or the proposed Maximal Entropy loss is capable of outperforming the original deep features in many cases. As detailed in the literature comparison, part of the adopted negative images clearly boosted the performance of our method whereas others encompassing distinct domains as well as different data distribution were not adequate and contributed little to the approach accuracy. How to obtain or generate more effective negative samples will be investigated in future work.

Acknowledgments

We are thankful to the Brazilian National Council for Scientific and Technological Development – CNPq (Grants 309953/2019-7 and 203402/2020-0), the Minas Gerais Research Foundation – FAPEMIG (Grant PPM-00540-17), the Federal University of Minas Gerais (UFMG) and, especially, the University of Zürich (UZH).

References

- [1] Kashmir Hill. Wrongfully accused by an algorithm. *New York Times*, 6 2020. <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>. 1
- [2] Manuel Günther, Akshay Raj Dhamija, and Terrance E. Boult. Watchlist adaptation: Protecting the innocent. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2020. 1, 2, 3, 5
- [3] Tony Romm. Amazon’s facial-recognition tool misidentified 28 lawmakers as people arrested for a crime, study finds. *Washington Post*, July 2017. 1
- [4] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision (ECCV)*. Springer, 2016. 2
- [5] Akshay Raj Dhamija, Manuel Günther, and Terrance E. Boult. Reducing network agnostophobia. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1, 4, 5, 8
- [6] Andres Palechor, Annesha Bhounik, and Manuel Günther. Large-scale open-set classification protocols for imagenet. In *Winter Conference on Applications of Computer Vision (WACV)*. CVF/IEEE, January 2023. 1
- [7] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, Univ. of Massachusetts, Amherst, 2007. 2, 3
- [8] Brianna Maze, Jocelyn Adams, James A. Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K. Jain, W. Tyler Niggel, Janet Anderson, Jordan Cheney, and Patrick Grother. IARPA Janus Benchmark - C: Face dataset and protocol. In *International Conference on Biometrics (ICB)*, 2018. 3
- [9] Manuel Günther, Peiyun Hu, Christian Herrmann, Chi-Ho Chan, Min Jiang, Shufan Yang, Akshay Raj Dhamija, Deva Ramanan, Jürgen Beyerer, Josef Kittler, et al. Unconstrained face detection and open-set face recognition challenge. In *International Joint Conference on Biometrics (IJCB)*. IEEE, 2017. 2, 3, 6
- [10] Chunchun Li, Manuel Günther, and Terrance E. Boult. ECLIPSE: Ensembles of centroids leveraging iteratively processed spatial eclipse clustering. In *Winter Conference on Applications of Computer Vision (WACV)*, 2018. 2, 3, 6
- [11] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. VGGFace2: A dataset for recognising faces across pose and age. In *Automatic Face & Gesture Recognition (FG)*. IEEE, 2018. 2, 3
- [12] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3, 6
- [13] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference (BMVC)*, 2015.
- [14] Jun-Cheng Chen, Vishal M Patel, and Rama Chellappa. Unconstrained face verification using deep CNN features. In *Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [15] 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/CVF conference on computer vision and pattern recognition*, 2017. 2
- [16] 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/CVF conference on computer vision and pattern recognition*, 2018.
- [17] Yifan Sun, Changmao Cheng, Yuhan Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 2

- [18] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Mag-Face: A universal representation for face recognition and quality assessment. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [19] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [20] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [21] Swami Sankaranarayanan, Azadeh Alavi, Carlos D. Castillo, and Rama Chellappa. Triplet probabilistic embedding for face verification and clustering. In *Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2016. 2
- [22] Rafael Henrique Vareto, Samira Silva, Filipe Costa, and William Robson Schwartz. Towards open-set face recognition using hashing functions. In *International Joint Conference on Biometrics (IJCB)*, 2017. 2
- [23] Aruni Roy Chowdhury, Tsung-Yu Lin, Subhansu Maji, and Erik Learned-Miller. One-to-many face recognition with bilinear CNNs. In *Winter Conference on Applications of Computer Vision (WACV)*, 2016. 2
- [24] Cassio Elias Dos Santos and William Robson Schwartz. Extending face identification to open-set face recognition. In *Brazilian Symposium on Computer Graphics and Image Processing (SIB-GRAPI)*, 2014. 2
- [25] James Henrydoss, Steve Cruz, Chunchun Li, Manuel Günther, and Terrance E. Boult. Enhancing open-set recognition using clustering-based extreme value machine (C-EVM). In *International Conference on Big Data (BigData)*. IEEE, 2020. 2
- [26] Rafael Henrique Vareto and William Robson Schwartz. Unconstrained face identification using ensembles trained on clustered data. In *International Joint Conference on Biometrics (IJCB)*, 2020. 2, 6
- [27] Mehadi Hassen and Philip K Chan. Learning a neural-network-based representation for open set recognition. In *International Conference on Data Mining*. SIAM, 2020. 2
- [28] Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Learning placeholders for open-set recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [29] Pramuditha Perera, Vlad I Morariu, Rajiv Jain, Varun Manjunatha, Curtis Wigington, Vicente Ordonez, and Vishal M. Patel. Generative-discriminative feature representations for open-set recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [30] Shu Kong and Deva Ramanan. Opengan: Open-set recognition via open data generation. In *International Conference on Computer Vision (ICCV)*, 2021.
- [31] Zhongqi Yue, Tan Wang, Qianru Sun, Xian-Sheng Hua, and Hanwang Zhang. Counterfactual zero-shot and open-set visual recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [32] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar datasets-10. URL <http://www.cs.toronto.edu/kriz/cifar.html>, 2010. 3
- [33] Yann LeCun, Corinna Cortes, and Christopher J.C. Burges. Mnist handwritten digits dataset. URL <http://yann.lecun.com/exdb/mnist/index.html>, 1998. 3
- [34] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *Advances in Neural Information Processing Systems (NeurIPS) Workshop*, 2011. 3
- [35] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2009. 3
- [36] Paul R Cohen and David Jensen. Overfitting explained. In *International Workshop on Artificial Intelligence and Statistics*. PMLR, 1997. 3
- [37] Manuel Günther, Steve Cruz, Ethan M. Rudd, and Terrance E. Boult. Toward open-set face recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. 3, 6, 7
- [38] Yoanna Martindez-Diaz, Luis S Luevano, Heydi Mendez-Vazquez, Miguel Nicolas-Diaz, Leonardo Chang, and Miguel Gonzalez-Mendoza. Shufflefacenet: A lightweight face architecture for efficient and highly-accurate face recognition. In *International Conference on Computer Vision (ICCV) Workshops*, 2019. 3
- [39] Archana Sapkota and Terrance E. Boult. Large scale unconstrained open set face database. In *Biometrics Theory, Applications and Systems (BTAS)*, 2013. 3
- [40] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *International Conference on Machine Learning*, pages 507–516. PMLR, 2016. 4
- [41] Xuezhi Liang, Xiaobo Wang, Zhen Lei, Shengcai Liao, and Stan Z Li. Soft-margin softmax for deep classification. In *International Conference on Neural Information Processing (ICONIP)*. Springer, 2017. 4
- [42] P. Jonathon Phillips, Patrick Grother, and Ross Micheals. *Handbook of Face Recognition*, chapter Evaluation Methods in Face Recognition. Springer, 2nd edition, 2011. 5
- [43] N Poh, C Chan, J Kittler, J Fierrez, and J Galbally. Beat-biometrics evaluation and testing: Description of metrics for the evaluation of biometric performance. Technical report, IDIAP, 2012.
- [44] Patrick Grother, Mei Ngan, and Kayee Hanaoka. Face recognition vendor test (frvt) part 2: Identification. Technical report, National Institute of Standards and Technology, 2022. 5
- [45] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *Signal Processing Letters*, 23(10), 2016. 6
- [46] André Anjos, Laurent El-Shafey, Roy Wallace, Manuel Günther, Christopher McCool, and Sébastien Marcel. Bob: a free signal processing and machine learning toolbox for researchers. In *ACM International Conference on Multimedia (ACMMM)*, pages 1449–1452, 2012. 6
- [47] Manuel Günther, Roy Wallace, and Sébastien Marcel. An open source framework for standardized comparisons of face recognition algorithms. In *European Conference on Computer Vision (ECCV) Workshops and Demonstrations*, pages 547–556. Springer, 2012. 6