

[Unconstrained] Open-set Face Recognition using Ensembles trained on Clustered Data

Rafael Henrique Vareto and William Robson Schwartz
Smart Sense Laboratory, Department of Computer Science
Universidade Federal de Minas Gerais, Belo Horizonte, Brazil
{rafaelvareto, william}@dcc.ufmg.br

Abstract

Open-set face recognition describes a scenario where unknown subjects, unseen during training stage, appear on test time. Not only it requires methods that accurately identify individuals of interest, but also demands approaches that effectively deal with unfamiliar faces. This work details a scalable open-set face identification approach to galleries composed of hundreds and thousands of subjects. It is composed of clustering and ensemble of binary learning algorithms that estimates when query face samples belong to the face gallery and then retrieves their correct identity. The approach selects the most suitable gallery subjects and use the ensemble to improve prediction performance. We carry out experiments on well-known LFW and YTF benchmarks. Results show that competitive performance can be achieved even when targeting scalability.

1. Introduction

Biometrics is used to identify and authenticate individuals using a set of discriminative physical and behavioral characteristics, which is inherent to every person. Few biometric traits are simultaneously non-invasive and accurate like face recognition. In fact, face recognition has been in use for decades and persists as one of the most widely used biometric traits.

Most face recognition approaches analyze facial geometries and their similarity to one another as they attempt to correctly identify a subject against a list of previously registered individuals (face gallery), containing faces with similar characteristics. However, face recognition is a broad term, commonly used to indicate three closely-related sub-tasks [1]: *face verification*, which focuses on determining whether a pair of different face images corresponds to the same subject; *closed-set face identification*, a task that presents a query image against a number of previously cataloged faces, ensuring that the probe face always contains a corresponding identity in the gallery set; and *open-set face identification*, which is similar to closed-set identification with the distinction that it does not guarantee that all query subjects are registered in the face gallery.

Both closed and open-set face identification tasks can be described as a $1 : N$ matching problem, where N equals the number of subjects enrolled in the gallery set. The former has training and testing data drawn from uniform label and feature spaces whereas the latter comes across situations in which unseen individuals emerge unexpectedly. More precisely, along with finding out which identity from the face gallery best matches an unknown face sample, open-set face identification systems first have to check whether the probe image indeed belongs to any of the registered subjects. In summary, open-set approaches behave like closed-set identification for known individuals and also label non-identified persons with an “unknown subject” category.

There is a great demand for open-set technology as the vast majority of real-world identification problems consist of a finite number of persons of interest in comparison to innumerable unknown individuals. As a clear illustration, think of an identification application for law enforcement agencies where lawbreakers’ identities are doubtless of interest; on the other hand, an infinite number of law-abiding citizens are not of concern. For that reason, open-set algorithms should dismiss all unwanted subjects and focus on identifying potential suspects only. Still, most researchers have left open-set problems aside and channeled their efforts into closed-set face identification problems despite the real-world appeal [2–6]. Thusly, there is yet a lot to improve when it comes to open-set face identification.

The proposed approach incorporates efficient and straightforward techniques, namely Affinity Propagation Clustering (APC) algorithm [7] and an ensemble of Partial Least Squares (PLS) [8] models. We follow Breiman *et al.*’s [9] statement that the building of multiple learning algorithms, trained on randomly generated training sets, tends to achieve better predictive performance than the composing classifier alone. This approach is inspired by the work of Vareto *et al.* [10], which aggregates each model’s response as it sets up a vote-list histogram where each bin comprises a gallery-registered subject. But differently, for a given query image, the proposed method performs affinity propagation clustering of all gallery subjects and elects the most suitable clusters to train a binary ensemble of classifiers on prominent subjects’ face samples only.

We believe that vote-list histograms react in a distinctive manner when query face samples correspond to identities previously registered in the gallery set, but present a uniform behavior when probe face images have no corresponding match. In fact, we hypothesize that when a probe sample is known, most classifiers would vote for the correct identity or otherwise distribute the votes among distinct gallery-registered individuals.

Unlike *one-against-all* learning schemes where generally unbalanced classification models are learned for each person of interest [11], the proposed approach is not directly dependent on the number of known subjects since clustering the gallery individuals and selecting the top k clusters significantly reduce the computational time. Actually, the method consists of balanced data splits. Qu *et al.* [12] show that unbalanced class distributions may cause conventional classifiers, such as support vector machines [13], to treat a few-sample class like noises and push the classification boundaries so that it may benefit the majority class, culminating in a precision drop of the outnumbered class. Providentially, the way the proposed method splits data guarantees symmetrical division between positive and negative classes for every model belonging to the ensemble.

To address complexity issues intrinsic to open-set recognition problems, we follow different literature protocols designed for open-set face evaluation on two well-known datasets: Labeled Faces in the Wild (LFW) [14] and Youtube Faces (YTF) [15]. We also check whether associating clustering and learning algorithms as an ensemble is an adequate form of outperforming more complex models. Then, we evaluate how the method’s performance responds to variable ensemble sizes and investigate whether it is possible to offer a trade-off between accuracy and simplicity. According to experimental results, our approach reports competitive matching accuracy in comparison with other state-of-the-art works on LFW and YTF databases.

The main contributions of this work are: (i) a scalable open-set approach, dismissing the need of retraining with the addition of new individuals; (ii) an easy-to-implement method with few trade-off parameters to be estimated; (iii) an algorithm that reduces the training space and learns small ensembles to improve prediction performance; (iv) a comprehensive discussion on the carried-out experiments, considering different protocols; (v) the development of a baseline that can be considered in future works regarding open-set face recognition.

2. Related Works

Contemporary years have witnessed a significant development on the face recognition field [16–18]. Although IBM stated approximately 50 years ago that humans could be recognized at a computer terminal by something they carry or by a personal characteristic [19], only recently has

open-set recognition been explored in the literature. On the other hand, there have been more works intended to solve closed-set and verification problems, either in unrestrained scenarios or in relatively small datasets [20–22].

Li *et al.* [23] developed a transductive inference that includes an option for rejecting unknown subjects. Kamgar *et al.* [24] identified decision regions in the high-dimensional face space by generating two large sets of borderline images through morphing techniques. Santos *et al.* [25] came up with five approaches, such as commonplace attributes among registered subjects, margin separation or distribution patterns between identification responses. Vareto *et al.* [10] proposed HPLS, assuming that vote-list histograms behave differently for queries comprising gallery-enrolled subjects and unknown individuals as histograms present highlighted bins when the query image matches an identity of interest. Yu *et al.* [26] designed a new loss function that utilizes unlabeled data to further enhance the discriminativeness of the learned feature representation.

The works mentioned above established different open-set protocols for a variety of datasets. This forges a problematic situation where it is almost impossible to make legitimate comparisons of research works since they do not follow a mainstream experimental evaluation. Some of these methods also overlook the identification stage as they only estimate whether a query sample belongs to the gallery set, without returning its actual identity. All these limitations raise the following question: “What is the most accurate algorithm for open-set face problems under certain shared conditions?”. As a result, the following literature works adopted either LFW [14] or YTF [15] datasets and devised new protocols for open-set assessment.

Best-Rowden *et al.* [17] boosted the likelihood of identifying subjects by fusing several media types. Liao *et al.* [27] evaluated seven learning algorithms taking as input three different feature descriptors. Sun *et al.* [28] designed two deep neural networks with prefix DEEPID to learn a richer pool of facial features. Martinez *et al.* [29] propose a lightweight architecture, SHUFFLEFACE, that use less than 4.5 million parameters. Günther *et al.* [30] introduced *known unknowns* (uninteresting persons) and *unknown unknowns* (people never seen before) individuals as they compare three similarity-assessing algorithms in deep feature spaces: Cosine similarity (COS), Linear Discriminant Analysis (LDA), and Extreme Value Machine (EVM) [31].

Despite the fact that the aforementioned approaches attained significant progress in the test-subject search, they still face scalability problems in cases of restricted training time or excessive data. The low computational time performance is probably due to the fact that these approaches still present linear asymptotic complexity with the gallery size. In such situations, unsatisfactory performance is observed when there are numerous individuals in the face gallery.

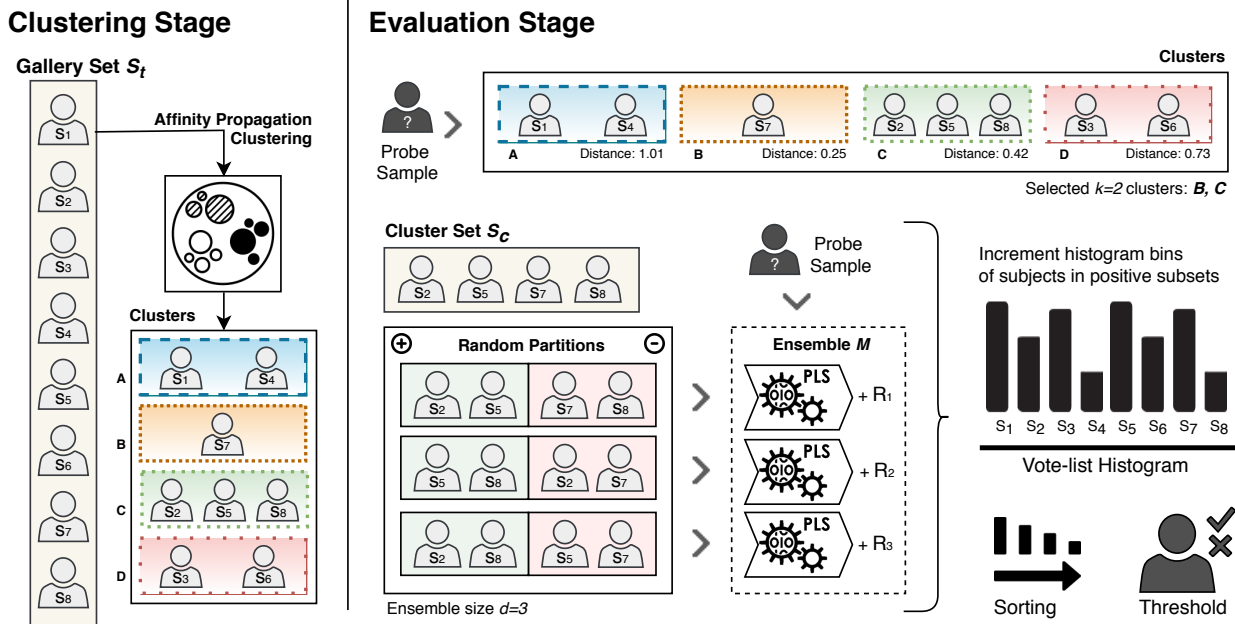


Figure 1: A pipeline illustration of proposed approach having an ensemble of classifiers with size $d = 3$ and selecting best $k = 2$ clusters. **Clustering:** Training data S_t is initially partitioned into APC-generated subgroups. Then, when a query sample q is presented to the collection of clusters, APC algorithm picks the most similar clusters (B and C) and sets up a training subset S_c . **Online Training:** Training subjects from S_c are randomly partitioned into d positive and negative subsets. Random partitions feed PLS models in the interest of training the ensemble M . **Testing:** The very same query sample q is presented to the ensemble in pursuance of response values $r_i \in R$. These values are used to increment the histogram bins of individuals $s_j \in S$ randomly sampled in the positive subset (green box). The histogram is sorted in descending order and becomes a ranking. The thresholding is employed to find out whether the probe sample belongs to the face gallery.

Few literature datasets introduce open-set face identification protocols. In fact, IJB-A [32] seems to be the first open-set benchmark containing “in-the-wild” image and video samples. Recently, Maze *et al.* [33] designed IJB-C on the lookout for evaluating end-to-end systems, that is, algorithms that jointly detect faces and recognize people. IJB-C includes a total of 31,334 images and 117,542 frames from 11,779 videos. The three benchmarks are produced by IARPA’s Janus program, which in 2019, decided to no longer distribute the datasets for an undisclosed period. Similarly, UCCS’s open-set surveillance dataset and respective challenge [34] are temporarily suspended due to metadata rectification for the third competition edition.

Recent open-set works dismiss IARPA’s datasets as some investigators claim they contain missing data annotations, low-quality images and videos [28, 30]. Günther *et al.* [30] declare that several issues must be worked out before researchers are able to correctly address the open-set problem. Some investigators have adopted LFW/YTF as the “official” benchmark for open-set face tasks due to the novel protocols released in subsequent research works [27, 30, 35]. They intend to come up against the lack of consensus of open-set biometrics. In the experiments, we evaluate the proposed method on LFW and YTF datasets as we adjust it to meet already-explored open-set protocols.

3. Proposed Method

Noisy training sets may mislead a method’s identification and the adoption of subsets of big data is a manner of removing meaningless information. Clustering reduces the computation time as it avoids comparing a probe image to all training classes, especially when a gallery set is large. The designed method, illustrated in Figure 1, is made of the Affinity Propagation Clustering (APC) [7] algorithm and an ensemble of binary Partial Least Squares (PLS) [8].

Gallery-composing subjects are submitted to a clustering algorithm, building the collection of clusters C . Given a query face sample q , individuals from top k clusters that most resemble q are singled out and their corresponding image samples are employed to learn d classifiers $m \in M^d$. Then, probe q is projected onto the embedding in favor of determining whether its identity lies in the gallery set or comprise an unknown subject. In fact, probe sample q ’s corresponding feature descriptor is presented to each model m in search of its response value r . Each value $r_i \in R^d$ is added to a vote-list histogram of size $|S_t|$ as it increments the bins of those subjects in the i -th positive subset of model $m_i \in M^d$. If the algorithm establishes that a probe image corresponds to an enrolled identity, the vote-list histogram turns into a list of candidates [10].

3.1. Affinity Propagation Clustering

Affinity Propagation is a centroid-based clustering algorithm that automatically chooses the number of clusters. Not only does it group facial data samples into different clusters, but it also estimates a representative sample (exemplar) for each one of them. In our context, it takes in a collection of real-valued similarities between face samples, where the similarity function $p(s_i, s_j)$ indicates how well a face image from subject s_j is suitable to be an exemplar for subject s_i . APC takes as input facial features and identifies exemplars based on certain norms, such as instances similarity and availability. It iteratively searches for clustering arrangements that satisfy its optimization criterion and stops when cluster boundaries stand consistent over a number of iterations.

Clusters Selection. The subjects $s \in S_t$ available in the gallery set input the APC procedure, resulting in the collection of clusters C . Whenever a probe face sample q turns up, the proposed approach projects q onto the clustered data $c \in C$, comparing it to the estimated representative exemplars (centroids). The first k clusters from collection C , containing individuals that most resemble query sample q , are maintained for ensemble learning as they constitute a smaller training set called S_c .

Container S_c contains a list of subject classes that favors high *interclass* similarity since the clustering strategy facilitates the picking of “look-alike” individuals. It eliminates most undesirable subjects and, consequently, endorses the learning of more robust classifiers as it involves distinguishing among visually-related samples. According to Caruana *et al.* [36], the building of more discriminating models allows the method to be optimized with respect to performance metrics such as accuracy, cross entropy and mean precision to mention a few.

3.2. Partial Least Squares Ensemble

The proposed ensemble is an adaptation of a technique called Bootstrap Aggregating [9], well-known for the acronym *bagging*. Several machine learning algorithms, like CNN, face instability issues when dealing with unbalanced training data. Favorable results can be attained when bagging is employed since the adoption of multiple learning algorithms tends to accomplish better predictive accuracy than standalone regression and classification models.

Online Training. Most bagging techniques consist of uniformly sampling from a face dataset with replacement, generating smaller training subsets in which some subject samples may be repeated. However, the developed method employs d random fifty-fifty disjoint splits: each random partitioning encompasses all individuals available in cluster set S_c by allocating half of them in the positive and the other half in the negative subset. The algorithm guarantees a com-

plete utilization of container S_c , ensures balanced data division, and implicitly provides each known subject with a binary string of dimension d .

The d random partitions provide a mapping function that attributes all subjects $s \in S_c$ to a Hamming embedding $h(s) \in \{\pm 1\}^d$ where the similarity between two strings indicates the Hamming distance between their images [37]. The chances of any two individuals sharing the very same binary string decreases as the number of regression models d expands. Note that the algorithm learns d different PLS models assigning feature descriptors obtained from individuals in the positive subset with target value $+1$ whereas those in the negative subset are assigned to target value -1 . That is, the assignment of a subject s to the i -th positive or negative subset, which is associated with model $m_i \in M^d$, designates the i -th bit in the Hamming embedding. Each model $m \in M^d$ contains a record of the individuals randomly included in its corresponding positive subset.

Testing. It is worthwhile to mention that the approach is equivalent to estimating Hamming embeddings for all subjects in the gallery set and then comparing them with the binary string generated for query sample q . Given q , the algorithm initializes the vote-list histogram and sets all bins (each bin corresponds to a subject $s \in S_t$) with zero values.

When q 's feature descriptor is presented to regression models $m \in M^d$, they return an array R^d composed of d response values r ranging between -1 and $+1$. Each model $m_i \in M$ holds its own list of subjects from the gallery set assigned to the i -th positive subset. More precisely, only histogram bins of individuals enlisted in the i -th positive subset are incremented by r_i . In addition, response values $r_i \in R$ go through a filtering function $r_i = \max(0, r_i)$, which resembles the ReLU activation function.

We expect that when a *known* probe sample is projected onto the ensemble, the regression response $r_i \in R^d$ associated with regression model $m_i \in M^d$ will indicate whether its corresponding class belongs to the i -th positive or negative subset. Particularly, it implies that a probe sample q resembles subjects from the i -th positive subset if r_i converges towards $+1$ in the testing stage. If the regression response r_i is negative, it suggests that the query sample q belongs to someone enrolled in the i -th negative subset. Then, the algorithm performs no increment in the vote-list histogram since negative responses are filtered out by the ReLU-like function.

Gallery Detection. As a query sample q is projected to all regression models $m \in M$ and the vote-list bins are progressively incremented, the proposed method sorts the histogram in descending order in behalf of creating a ranking of candidates. After the ranking generation, the vote-list histogram is arranged in such a way that individuals with higher probability of matching probe q lie on top of the candidates ranking.

Since we are dealing with an open-set task, it is essential to find out if the gallery set indeed contains an identity that matches q . Therefore, the same list of candidates is explored as the approach computes the ratio of the top scorer to the mean of the next two following individuals, candidates two and three [10]. If the ratio is higher than a specified threshold, the method considers the probe sample as a known subject. If not, the pipeline halts since there is no further reason to continue searching for a subject holding low probability of being previously enrolled in the gallery.

3.3. Technical Analysis

A practical advantage of the proposed algorithm is that the search for faces resembling a given probe image is reduced to a search in the metric embedding space, which is executed quickly. The adoption of the vote-list histogram is a manner of weighting the hamming bits when individuals lie in positive subsets. With clustering, given a dataset of numerous faces and a probe image, the idea is to retrieve identities that are similar to the query without comparing it to every subject (class) enrolled in the gallery set.

The method’s underlying structure makes the ensemble size independent from the number of gallery-registered subjects. Besides, no retraining is required when extra individuals are inserted into the training set. With this autonomous characteristic, the combination of APC clusters and PLS ensembles is capable of achieving satisfactory results and outperform literature techniques by enclosing far fewer models than traditional multi-class classification schemes [11]. Consequently, the designed algorithm provides scalability to galleries composed of numerous individuals and on which regular open-set face recognition methods may probably fail to respond in low computational time [38].

4. Experimental Results

This section presents the experimental evaluation of the approach described in Section 3. From now on, we refer to the developed method as COT since it consists of a clustering strategy followed by training a collection of PLS models.

4.1. Setup

We explore the Scikit-Learn library for Python, an efficient open-source tool for data analysis and mining. COT operates on high-quality deep features, extracted with publicly available VGGFACE network [39], a convolutional neural network designed for face detection and recognition, based on VGG16, RESNET50 and SENET50 architectures. TensorFlow is the adopted neural network library, high-leveled with Keras API for fast experimentation. All algorithm evaluations are performed on Intel Xeon E5-2630 CPU with 2,30 GHz and 12GB of RAM using Ubuntu 18.04 LTS operating system, no more than 8GB of RAM was required though.

4.2. Datasets

Experiments regarding parameter validation and comparison with state-of-the-art methods are conducted with unofficial open-set protocols on LFW and YTF benchmarks¹.

LFW [14]. A database of face photographs containing approximately 13,000 uncontrolled face images of almost six thousand individuals. The original database includes four different LFW sets images as well as three different types of aligned images. Following a recent work from literature [30], we adopt the funneled aligned LFW images.

YTF [15]. Contains face videos captured for investigating the problem of unconstrained face verification in videos. It consists of 3,425 videos of 1,595 different people. Organizers provide the entire dataset broken into frames and pre-processed with facial detection and alignment.

4.3. Protocols

In this experimental evaluation, we adopt conventions proposed in recent literature works, which are identified by each work’s corresponding first author’s name:

ROWDEN protocol [17]. This protocol partitions LFW so that the gallery set consists of 596 identities, in which each individual has at least two training images and exactly a single probe face sample. The remaining 4,494 face samples comprise the set of impostor probe images.

GÜNTHER protocol [30]. This protocol splits LFW subjects into three disjoint groups: 602 identities containing more than three face samples compose the *known* subset. Other 1,070 individuals holding two or three images constitute the *known unknowns* subset (distractors). Lastly, 4,096 identities with a single image each only available during test time.

MARTINEZ protocol [35]. This protocol randomly divides the YTF into ten training and test subsets. It is characterized by the concept of *openness* (OP): defined as the ratio between the genuine comparisons and the impostor comparisons in the probe set.

4.4. Metrics

The Receiver Operating Characteristic (ROC), its associated Area Under Curve (AUC), and the Cumulative Match Characteristic (CMC) are widely employed in closed-set tasks [41]. We also incorporate the Detection and Identification Rate (DIR) and False Alarm Rate (FAR). Plotting DIR vs. FAR produces a chart known as Open-set ROC, a metric generally used to evaluate approaches composed by filtering and identification steps [42].

¹FYI. It is important to mention that throughout the development process, we have requested datasets designed for open-set face tasks, such as IJB-A [32], IJB-C [33], MegaFace [40], and the UCCS competition [34]. We have contacted their administrators who stated that these benchmarks are either unavailable for copyright issues or erroneous metadata. Consequently, these face datasets are not included in our experimental section.

4.5. Parameter Selection

Experiments carried out in this subsection contemplate VGGFACE’s RESNET50 architecture and LFW dataset in a closed-set fashion. We follow both ROWDEN and GÜNTHER protocols to a certain extent: a random dataset partition, selecting 602 subjects to compose gallery and probe sets as well as setting the maximum of three samples for training and a single sample for testing.

PLS dimensions. PLS regression models only require a single parameter: the number of components in the latent space. Table 1 presents CMC results as we gradually increase the number of PLS components from 5 to 20, but keep fixed ensemble size d of 50 PLS models. Since there is a minor advantage in using a particular dimension size, we set the number of PLS components c to 10 for all following experiments after 30 executions.

PLS dimensions		5	7	10	12	14	20
Rank-01	AVG	93.0	94.4	96.7	96.2	94.3	93.5
	STD	2.02	1.67	1.62	1.66	1.46	1.40
Rank-05	AVG	94.8	96.0	98.1	98.3	96.4	95.7
	STD	1.25	1.06	1.03	0.99	1.10	1.15
Rank-10	AVG	95.8	96.2	98.9	98.9	97.8	96.7
	STD	1.07	1.05	0.93	1.00	0.83	0.86

Table 1: Variable number of PLS dimensions and their corresponding CMC averaged results (%) on different Ranks.

PLS models. Progressively adding regression models results in extra random partitions, further models training and extra probe feature projections during evaluation time. Initially, a small embedding size d does not seem adequate for discerning when an individual is registered in the gallery set and, consequently, retrieve the correct identity.

PLS models		30	40	50	60	80	100
Rank-01	AVG	93.0	95.4	96.7	97.2	97.3	97.6
	STD	2.40	1.64	1.62	1.15	1.14	1.03
Rank-05	AVG	95.3	97.3	98.1	98.6	98.9	99.0
	STD	1.99	1.23	1.03	0.66	0.43	0.35
Rank-10	AVG	97.9	98.6	98.9	99.1	99.2	99.2
	STD	1.38	1.12	0.93	0.42	0.29	0.12

Table 2: Variable number of PLS models and their corresponding CMC averaged results (%) on different Ranks.

Table 2 exposes how the number of PLS models affects the implemented method. There is a meaningful improvement when varying the ensemble size from 30 to 100. However, we observe no significant accuracy increase when the number of regression models exceeds 50 learning algorithms. This experiment advocates the claim in which we declare that the ensemble size is not directly dependent on the number of known individuals. Therefore, we fix the number of PLS models on 60 in the subsequent ex-

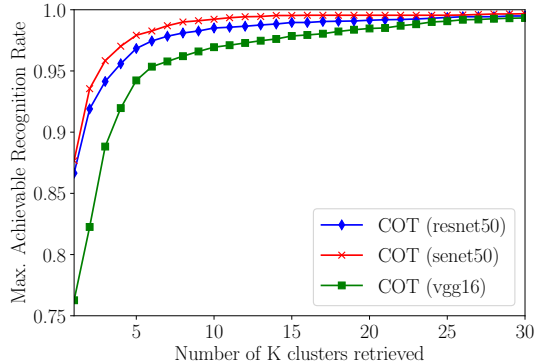


Figure 2: Comparison of the number of clusters retrieved and the maximum achievable recognition rate.

periments. We are inclined to believe that either a reduced number of samples per identity or an expanded number of individuals enrolled in the gallery set during training time may require additional PLS regression models in order to keep CMC Rank-1 values high.

Potential Accuracy. The Maximum Achievable Recognition Rate (MARR) [38] permits the evaluation of the filtering approach performance, adopted to reduce the PLS ensemble size. It demonstrates how an error originated during the picking of best clusters propagates to the rest of the pipeline.

MARR assumes that a flawless face identification method is employed and, therefore, indicates the upper bound rate for both filtering and identification pipeline. Failing to select the best k clusters results in trouble returning the correct identity for a given probe image. Figure 2 demonstrates that an insufficient number of retrieved clusters severely diminishes the approach’s capacity. However, the ensemble may be able to fix the inferiority of VGG16 with larger numbers of k . Based on the experiments, we set k to 20 in all experiments described in the following section.

4.6. Literature Comparison

On the contrary of previous experiments that only deal with known individuals, this subsection focuses on assessing the complete open-set identification pipeline by contrasting the proposed COT approach with state-of-the-art works. Besides, we incorporate both VGGFACE’s architectures: RESNET50 and SENET50.

Figure 3 presents the comparison between literature methods in total accordance with GÜNTHER protocol. Note that VGGFACE architectures embedded in COT share similar performance nature. The CMC chart demonstrates a superior performance of COT up to Rank-10. The proposed method outperformed all literature works when evaluating the open-set ROC curve, displayed in the chart on the right. The DIR vs. FAR plot evaluates both types of unknowns, which consists of the average evaluated methods’ performance on known unknowns and unknown unknowns.

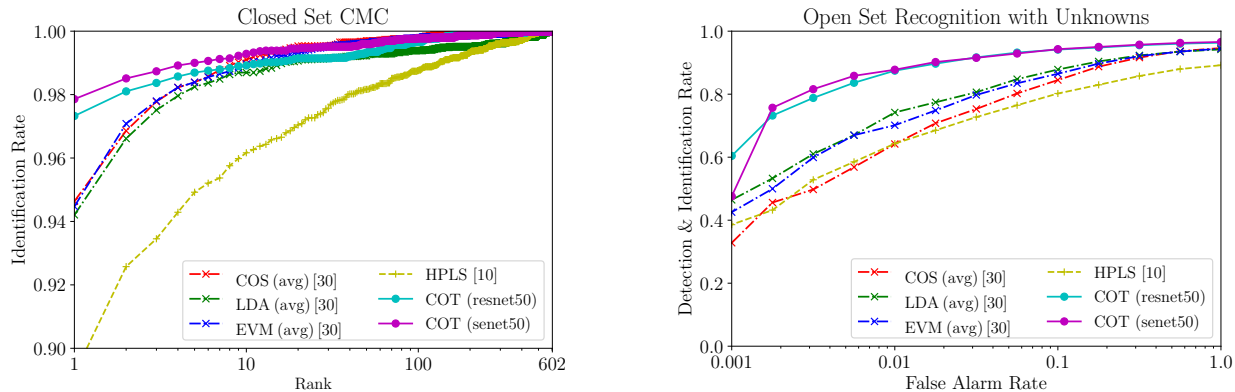


Figure 3: CMC and open-set DIR comparison between methods following GÜNTHER protocol on LFW.

Method	DEEPIID3	DEEPIID2	HPLS	COT _{RES50}
Rank-1	96.0	95.0	90.7	98.1
DIR@0.01	81.4	80.7	64.3	82.3

Table 3: Comparison between methods following ROWDEN protocol on LFW dataset.

In Figure 3, approximately 88% of known subjects are correctly identified with our proposed method. An open-set ROC curves’s top left corner indicates the optimal point. For a clear understanding of the results, a false alarm rate of 0.01 implies that one out of 100 unknown individuals are mistakenly assigned to an identity from the face gallery². Table 3 compares closed-set *Rank-1* and open-set DIR identification rates when the number of false alarms (FAR) is fixed on 1% under ROWDEN protocol. COT_{RESNET50} algorithm achieved best CMC@*Rank-1* accuracy of 98.1 as well as a DIR@FAR=0.01 of 82.3. We expected a better performance of the proposed method in this ROWDEN protocol as it provides more samples per class in the training set.

As mentioned in Sections 1 and 3, COT is not directly dependent on the number of classes, but to the number of subjects encompassed by the top k clusters. Both LFW protocols provide over 590 subject identities during training time yet the proposed method did not require an ensemble with more than 60 PLS regression models. In addition, since the random partitions generate balanced data splits, we can replace PLS by other long-established learning algorithms, such as support vector machines [13] and neural networks, without facing instability issues that may cause accuracy drop due to disproportionate training data.

Table 4 presents closed-set CMC results on the YTF database following MARTINEZ closed-set protocol (video-to-video). The proposed COT_{RESNET50} algorithm achieved best CMC@*Rank-1* accuracy of 87.65 when openness (OP)

²As an illustration, think of a biometric surveillance system at a football stadium that takes around 100 face pictures in a minute. The system is going to trigger one false alarm every sixty seconds as it searches for each one of them in a mugshot database. On the other hand, COT would successfully recognize nearly 90% of lawbreakers and fugitives.

Method	SHUFFLENET	SHUFFLEFACE	COT _{RES50}
OP=0.2	82.65	86.83	87.65
OP=0.5	79.22	85.52	83.79
OP=0.9	78.03	84.61	79.54

Table 4: Comparison between methods on YTF dataset.

was set to 0.2. However, as the gallery set increased, SHUFFLEFACE [29] obtained the best results. In summary, the proposed method’s generalization capability allows it to obtain good results in closed-set problems even though it was initially designed for open-set face tasks.

5. Conclusion

This work details a scalable open-set face identification approach to galleries with hundreds and thousands of subjects. The proposed algorithm combines a clustering algorithm with an ensemble of regression models, so that it either retrieves individuals from the gallery set that present substantial similarity to the probe image or considers them unknown subjects.

We evaluated different parameters to check how they would impact the proposed algorithm. All experiments were conducted on LFW and YTF, face databases initially proposed for face verification. We adopted third-party protocols with the intent of establishing appropriate comparisons with state-of-the-art methods. Both databases do not provide enough number of individuals for a pertinent study regarding scalability towards large face galleries. In summary, open-set face identification remains a difficult problem and more effort is necessary to make it fully effective.

Acknowledgments. The authors would like to thank the National Council for Scientific and Technological Development – CNPq (Grants 438629/2018-3 and 309953/2019-7), the Minas Gerais Research Foundation – FAPEMIG (Grants APQ-00567-14 and PPM-00540-17), the Coordination for the Improvement of Higher Education Personnel – CAPES (DeepEyes Project). This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

References

- [1] Rama Chellappa, Pawan Sinha, and P Jonathon Phillips. Face recognition by computers and humans. *Computer*, 2010. 1
- [2] Brendan F Klare and Anil K Jain. Heterogeneous face recognition using kernel prototype similarities. *PAMI*, 2012. 1
- [3] Dong Yi, Zhen Lei, and Stan Z Li. Towards pose robust face recognition. In *CVPR*. IEEE, 2013. 1
- [4] Meina Kan, Shiguang Shan, Hong Chang, and Xilin Chen. Stacked progressive auto-encoders (spae) for face recognition across poses. In *CVPR*, pages 1883–1890. IEEE, 2014. 1
- [5] Dong Yi, Zhen Lei, and Stan Z Li. Shared representation learning for heterogenous face recognition. In *FG*, pages 1–7. IEEE, 2015. 1
- [6] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *CVPR*. IEEE, 2017. 1
- [7] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *Science*, 2007. 1, 3
- [8] Roman Rosipal and Nicole Krämer. Overview and recent advances in partial least squares. In *ISOPWS*. Springer, 2005. 1, 3
- [9] Leo Breiman et al. Heuristics of instability and stabilization in model selection. *The annals of statistics*, 1996. 1, 4
- [10] Rafael Vareto, Samira Silva, Filipe Costa, and William Robson Schwartz. Towards open-set face recognition using hashing functions. In *IJCB*. IEEE, 2017. 1, 2, 3, 5
- [11] Guang-Bin Huang, Hongming Zhou, Xiaojuan Ding, and Rui Zhang. Extreme learning machine for regression and multiclass classification. *TSMC-B*, 2011. 2, 5
- [12] Hai-Ni Qu, Guo-Zheng Li, and Wei-Sheng Xu. An asymmetric classifier based on partial least squares. *PR*, 2010. 2
- [13] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008. 2, 7
- [14] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 2, 5
- [15] L Wolf, T Hassner, and I Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR*. IEEE, 2011. 2, 5
- [16] John Wright, Allen Y Yang, Arvind Ganesh, S Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *PAMI*, 2009. 2
- [17] Lacey Best-Rowden, Hu Han, Charles Otto, Brendan F Klare, and Anil K Jain. Unconstrained face recognition: Identifying a person of interest from a media collection. *TIFS*, 2014. 2, 5
- [18] Nate Crosswhite, Jeffrey Byrne, Chris Stauffer, Omkar Parkhi, Qiong Cao, and Andrew Zisserman. Template adaptation for face verification and identification. *IVC*, 2018. 2
- [19] Anil K Jain, Patrick Flynn, and Arun A. Ross. *Handbook of Biometrics*. Springer, 2008. 2
- [20] Xiaoyang Tan and William Triggs. Enhanced local texture feature sets for face recognition under difficult lighting conditions. *TIP*, pages 1635–1650, 2010. 2
- [21] Xiangyu Zhu, Junjie Yan, Dong Yi, Zhen Lei, and Stan Z Li. Discriminative 3d morphable model fitting. In *FG*, volume 1, pages 1–8. IEEE, 2015. 2
- [22] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu. Additive margin softmax for face verification. *SPL*, 2018. 2
- [23] Fayin Li and Harry Wechsler. Open set face recognition using transduction. *PAMI*, 2005. 2
- [24] Behrooz Kamgar-Parsi, Wallace Lawson, and Behzad Kamgar-Parsi. Toward development of a face recognition system for watchlist surveillance. *PAMI*, 2011. 2
- [25] Cassio E dos Santos Jr and William Robson Schwartz. Extending face identification to open-set face recognition. In *SIBGRAPI*. IEEE, 2014. 2
- [26] Haiming Yu, Yin Fan, Keyu Chen, He Yan, Xiangju Lu, Junhui Liu, and Danming Xie. Unknown identity rejection loss: Utilizing unlabeled data for face recognition. In *ICCVW*. IEEE, 2019. 2
- [27] Shengcai Liao, Zhen Lei, Dong Yi, and Stan Z Li. A benchmark study of large-scale unconstrained face recognition. In *IJCB*, pages 1–8. IEEE, 2014. 2, 3
- [28] Yi Sun, Ding Liang, Xiaogang Wang, and Xiaoou Tang. Deepid3: Face recognition with very deep neural networks. *arXiv preprint arXiv:1502.00873*, 2015. 2, 3
- [29] Yoanna Martınez-Dıaz, Luis S Luevano, Heydi Mendez-Vazquez, Miguel Nicolas-Dıaz, Leonardo Chang, and Miguel Gonzalez-Mendoza. Shufflefacenet: a lightweight face architecture for efficient and highly-accurate face recognition. In *ICCVW*, 2019. 2, 7
- [30] Manuel Günther, Steve Cruz, Ethan M Rudd, and Terrance E Boulton. Toward open-set face recognition. In *CVPRW*, pages 71–80. IEEE, 2017. 2, 3, 5
- [31] Ethan M Rudd, Lalit P Jain, Walter J Scheirer, and Terrance E Boulton. The extreme value machine. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):762–768, 2017. 2
- [32] Brendan F Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, and Anil K Jain. Pushing the frontiers of unconstrained face detection and recognition: Iarpa janus benchmark a. In *CVPR*, pages 1931–1939. IEEE, 2015. 3, 5
- [33] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *ICB*. IEEE, 2018. 3, 5
- [34] 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 *IJCB*. IEEE, 2017. 3, 5
- [35] Yoanna Martínez-Díaz, Heydi Méndez-Vázquez, Leyanis López-Avila, Leonardo Chang, L Enrique Sucar, and Massimo Tistarelli. Toward more realistic face recognition evaluation protocols for the youtube faces database. In *CVPRW*, 2018. 3, 5
- [36] Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. Ensemble selection from libraries of models. In *ICML*. ACM, 2004. 4
- [37] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *ECCV*. Springer, 2008. 4
- [38] Cassio E dos Santos Jr, Ewa Kijak, Guillaume Gravier, and William Robson Schwartz. Partial least squares for face hashing. *Neurocomputing*, 2016. 5, 6
- [39] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *FG*. IEEE, 2018. 5
- [40] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *CVPR*, pages 4873–4882, 2016. 5
- [41] Anil Jain, Lin Hong, and Sharath Pankanti. Biometric identification. *ACM*, 2000. 5
- [42] Anil K Jain and Stan Z Li. *Handbook of face recognition*. Springer, 2011. 5