

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/321347568>

# Training Convolutional Neural Networks with Limited Training Data for Ear Recognition in the Wild

Article · November 2017

DOI: 10.1109/FG.2017.123

CITATIONS

32

READS

348

4 authors:



**Žiga Emeršič**

University of Ljubljana

39 PUBLICATIONS 404 CITATIONS

[SEE PROFILE](#)



**Dejan Štepec**

University of Ljubljana

12 PUBLICATIONS 101 CITATIONS

[SEE PROFILE](#)



**Vitomir Štruc**

University of Ljubljana

151 PUBLICATIONS 1,938 CITATIONS

[SEE PROFILE](#)



**Peter Peer**

University of Ljubljana

128 PUBLICATIONS 1,693 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Special Issue "Deep Image Semantic Segmentation and Recognition" [View project](#)



Model-based gait recognition [View project](#)

# Training Convolutional Neural Networks with Limited Training Data for Ear Recognition in the Wild

Žiga Emeršič<sup>1</sup>, Dejan Štepec<sup>1</sup>, Vitomir Štruc<sup>2</sup> and Peter Peer<sup>1</sup>

<sup>1</sup> Faculty of Computer and Information Science, University of Ljubljana, Slovenia

Email: {ziga.emersic, dejan.stepec, peter.peer}@fri.uni-lj.si

<sup>2</sup> Faculty of Electrical Engineering, University of Ljubljana, Slovenia

Email: vitomir.struc@fe.uni-lj.si

**Abstract**—Identity recognition from ear images is an active field of research within the biometric community. The ability to capture ear images from a distance and in a covert manner makes ear recognition technology an appealing choice for surveillance and security applications as well as related application domains. In contrast to other biometric modalities, where large datasets captured in uncontrolled settings are readily available, datasets of ear images are still limited in size and mostly of laboratory-like quality. As a consequence, ear recognition technology has not benefited yet from advances in deep learning and convolutional neural networks (CNNs) and is still lacking behind other modalities that experienced significant performance gains owing to deep recognition technology. In this paper we address this problem and aim at building a CNN-based ear recognition model. We explore different strategies towards model training with limited amounts of training data and show that by selecting an appropriate model architecture, using aggressive data augmentation and selective learning on existing (pre-trained) models, we are able to learn an effective CNN-based model using a little more than 1300 training images. The result of our work is the first CNN-based approach to ear recognition that is also made publicly available to the research community. With our model we are able to improve on the rank one recognition rate of the previous state-of-the-art by more than 25% on a challenging dataset of ear images captured from the web (a.k.a. in the wild).

## I. INTRODUCTION

Convolutional neural networks (CNNs) have recently demonstrated impressive performance on various computer vision tasks such as semantic image segmentation [1], object detection and recognition [2], [3], [4], image super-resolution [5], [6], [7] and alike. One of the key factors contributing to this development is the availability of extensive corpora of training data that allow for the design and especially training of deep CNNs. For problems, where training data is in abundance, CNNs have pushed the state-of-the-art to new unprecedented heights, whereas in other areas, where training data is scarce, the impact of deep learning and CNNs has been limited. One such area is automatic person recognition from ear images, which has immense potential in forensic, security and surveillance applications, but lacks the necessary large-scale datasets to make full use

of recent developments in deep learning and convolutional neural networks.

While ear recognition is getting more and more popular over recent years, the existing datasets in this area are still limited to a few thousand images with a few hundred identities and are typically of laboratory-like quality with constrained variability. This is in stark contrast to the field of face recognition (or image classification), where datasets used for CNN training are nowadays measured in tens of millions of images with tens of thousands of identities. Moreover, these datasets are commonly harvested from the web and are therefore generally considered to be representative of real-world settings.

In this paper we address the problem of training CNNs with limited training data and strive to develop an effective CNN-based model for ear recognition. Existing approaches to CNN training with small amounts of training data typically include *i*) metric-learning approaches, where training is performed with image pairs (or even triplets) instead of single images [8], [9], *ii*) data augmentation techniques that in addition to geometric and color perturbations of the existing training data also include the generation of synthetic data samples [10], [11], [12], [13], and *iii*) using existing CNNs (trained for related recognition problems) as so-called “black-box” feature extractors, on top of which additional classifiers are trained and used for recognition [14]. Here, we build on these approaches and successfully develop a CNN model for ear recognition by exploring different strategies to network training, i.e.:

- Model architecture selection: we investigate different CNN architectures with different parameter-space sizes. Specifically, we consider an AlexNet-like architecture [15], [16], the 16-layer VGG model architecture [17], and the more recent SqueezeNet architecture [18].
- Aggressive data augmentation: we examine various data augmentation techniques and try to train CNN models for ear recognition from scratch using a (web-harvested) dataset of unconstrained ear images.
- Selective model learning: we consider pre-trained CNN trained initially for face or image classification and then learn only parts of the model to reduce the number of parameters that need to be estimated with the available training data.

The result of our development work is a CNN-based model for ear recognition that is able to increase the Rank-1 recognition rate on our test data by more than 25% compared to the existing state-of-the-art.

In summary, we make the following contributions in this paper:

- We develop the first CNN-based model for ear recognition using a limited amount of training data and significantly improve on the existing state-of-the-art on a challenging dataset of ear images gathered from the web.
- We evaluate, compare and discuss different strategies to training CNN-based recognition models with limited training data and elaborate on what works and what does not.
- We present comparative experiments between our CNN-based model and 7 state-of-the-art descriptor-based techniques for ear recognition.

The rest of the paper is structured as follows. In Section II we briefly review the existing work on ear recognition. In Section III we present the background and motivation for our work, describe the CNN architectures and strategies we followed during model learning. In Section IV we outline the experimental setup, datasets, and report experimental results. We conclude the paper in Section V, respectively.

## II. RELATED WORK

The goal of this section is to provide the reader with a brief high-level overview of existing techniques in the field of ear recognition and CNN-based recognition models. For a comprehensive overview of the two areas, the reader is referred to in-depth surveys on these topics [10], [19].

**Ear recognition:** According to Emeršič et al. [19] existing techniques for ear recognition can conveniently be grouped into *i)* geometric, *ii)* holistic, *iii)* local, and *iv)* hybrid techniques.

Geometric techniques describe the geometric properties of ears or derive geometry-related statistics that can be used for recognition. Since only information related to the ear geometry is used, it is straightforward to devise methods that are invariant to geometric distortions, such as rotation, scaling or even small perspective changes. Examples of techniques from this group were presented in [19], [20].

Holistic techniques rely on the global ear appearance and exploit representations that encode the ear structure as whole. As the appearance of ears varies significantly with pose or illumination, care needs to be taken before computing holistic features from input images and normalization techniques need to be applied to correct for these changes prior to feature extraction. Examples of global techniques can be found in [21], [22], [23].

Local approaches extract features from spatially-confined areas of an image without leaning on the global information describing the overall structure of the ears. The extracted features do not necessarily correspond to structurally meaningful parts of the ear, but can in general represent any point in the image. We distinguish two types of local

techniques: techniques that first detect keypoint locations in the image and then compute descriptors for each of the detected keypoints [24] and techniques that compute local descriptors densely over the entire image with no regard to the image's structural characteristics. Examples of local techniques include [19], [25], [26], [27].

The last groups of techniques, so-called hybrid techniques, combine elements from other categories or use multiple representations to increase the ear recognition performance [19]. Techniques from this group offer superior performance compared to competing techniques, but often at the cost of higher computational complexity [28], [29], [30]. As suggested by recent surveys on ear recognition [19], [20], [31], hybrid techniques together with local descriptor-based methods represent the current state-of-the-art in this area.

**CNN-based recognition models:** Various CNN architectures have been developed and presented in the literature over recent years. One of the most well-known problems that highlighted the potential and power of CNN-based approaches was object classification within the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [32]. The AlexNet architecture introduced by Krizhevsky et al. [16] achieved an unprecedented performance on the ImageNet data and triggered a surge in the use and popularity of CNN-based models in computer vision [16]. The next milestone in terms of ILSVRC results marked the introduction of the 16-layer VGG [17] architecture which provided an additional boost to the recognition rates of ILSVRC. These two architectures are nowadays also used in other tasks for example as part of larger architectures such as Faster-RCNNs [33] or SegNet [34], [35]. A more recent architecture, called ResNet, was presented by He et al. [36]. The architecture introduced shortcut connections to CNNs and made it possible to reliably train deep networks with several hundreds or even thousands of network layers.

To the best of our knowledge, no CNN-based methods for ear recognition have yet been presented in the literature, with the exception of Galdamez et al. [37], where the authors tried to build separate CNN models for each subject. The main reason for the lack of work in the field of CNN-based ear recognition is in our opinion the absence of large-scale datasets and the difficulty of effective training of deep model with small amounts of data as also suggested in [19].

## III. METHODOLOGY

In this work, we are interested in CNN-based models applicable for the task of closed-set identification from ear images. Our goal is to train CNN-based models that are able to determine the correct class (or identity) of an input ear image from a closed-set of predefined target identities. Thus, each CNN model considered features a softmax layer that returns a class-membership distribution over all available target classes when an input ear images is fed to the model. In the remainder of this section we outline the main strategies we explore to train an effective CNN-based model for ear recognition based on an limited amount of training data.

## A. Learning strategies

To train a CNN-based model for ear recognition we investigate the following learning strategies.

**Model architecture selection:** Different CNN-based models contain different numbers of model-parameters that need to be determined during training. This fact imposes certain requirements on the amount of training data that need to be available for parameter learning. While it is reasonable to assume that lighter architectures (i.e., architectures with less parameters) require less data to be trained, the convergence of the back-propagation (learning) procedure also needs to be considered, since certain architectures may facilitate faster learning. To study the impact of the model architecture on the learning procedure, we consider three popular CNN configurations, i.e.,

- *AlexNet* [16], which is one of the first successful CNN architectures initially applied for the problem of object classification on the ImageNet dataset [32]. AlexNet helped to popularize the field of deep learning and still represents a successful architecture used for various vision tasks. The model contains around 58.96 million parameters and 500,000 neurons and consists of five convolutional layers, (some of which are followed by max-pooling layers), and three fully-connected layers. The fully connected layers are followed by a softmax classifier that outputs class-membership probabilities for each relevant class. In this work, we used an AlexNet-like architecture available with the Caffe library (also referred to as CaffeNet [15]), where additional normalization layers are used after the pooling operations.
- *VGG-16* [17] is a representative of so-called *very deep* CNN models. The main characteristic of the network is the use of several consecutive convolutional layers with small  $3 \times 3$  kernels (or filters) – which is also the smallest filter size capable of encoding directional information. These stacks of  $3 \times 3$  convolutional layers are able to capture the same information as the larger filter used with AlexNet, but require significantly less parameters that need to be estimated during training. The  $3 \times 3$  filter stacks are interspersed with max-pooling layers which reduce the dimensionality of the activation maps produced by the model layers. The convolutional part of the VGG model is followed by three fully-connected layers with 4,096, 4,096 and 1,000 channels respectively. Finally a soft-max layer is used at the top of the network. The VGG model contains around 134.94 million parameters.
- *SqueezeNet* [18] represents a special example of a lightweight residual network (or ResNets), where squeeze layers, i.e., layers consisting only of  $1 \times 1$  convolutions, are added to the network with the goal of reducing the model size and number of weights that need to be tuned during training. Furthermore, no fully-connected layers are present in the network, instead an average-pooling layer is used at the top of the network to produce the final representation that serves as the

input to the softmax classifier. The SqueezeNet architecture also contains residual connections whose purpose is to make the back-propagation-based learning more efficient. The architecture used for our experiments contains the following layers from input to output: a convolutional layer, a max-pooling layer, three so-called fire modules, a max-pooling layer, another four fire modules, a max-pooling layer, one additional fire module, a convolutional layer and finally the average-pooling layer. In total, SqueezeNet model used in our experiments contains around 821 thousand parameters before pruning.

**Full model learning:** Learning CNN-based models from scratch is a difficult task that requires large amounts of appropriately annotated training data. The model parameters are learned with some form of back-propagation algorithm that tries to minimize a loss function defined over the output of the model. Commonly, the lower network layers learn to respond to primitive visual features, such as edges or corners, which are then gradually combined into more complex structures in the higher model layers until some data representation is learned in the fully connected layers of the CNN-based model driven by the learning criterion.

Since most CNN architecture contain several millions of parameters that need to be learned during training, even large datasets featuring tens of thousands of images are usually not sufficient to facilitate successful training. Aggressive data augmentation is, therefore, a must with CNN-based models. When trying to learn a CNN-based model from scratch, researchers typically augment the available training data by producing data variations with, e.g., geometric transformations, color modifications, addition of noise, and more recently also by synthesizing samples of artificial identities, as, for example, described in [13].

When investigating the use of *full model learning* for the problem of ear recognition, we also rely on aggressive data augmentation and increase the amount of available training data by up to a factor of 100. Since data augmentation is also used with the selective training (described next), we present our data augmentation steps in a separate section below.

**Selective model learning:** CNN-based models have been successfully applied to a number of visual recognition problems as outlined in the introductory section. These models typically share some common characteristics, which are reflected in the CNN-based models, especially at the lower model layers. It is, therefore, reasonable to assume that CNN-based models trained, for example, for object recognition should also be able to produce useful representations of more specific object classes, such as ears, along the model layers. While these representations are likely not optimal for the problem of ear recognition, existing pretrained CNN-based models should nevertheless represent a suitable initial model configuration that may be exploited to reduce the need for large amounts of training data. As part of our *selective model learning* strategy, we therefore initialize each of the considered architectures with parameters learned on the ImageNet dataset and fine-tune certain layers only. The

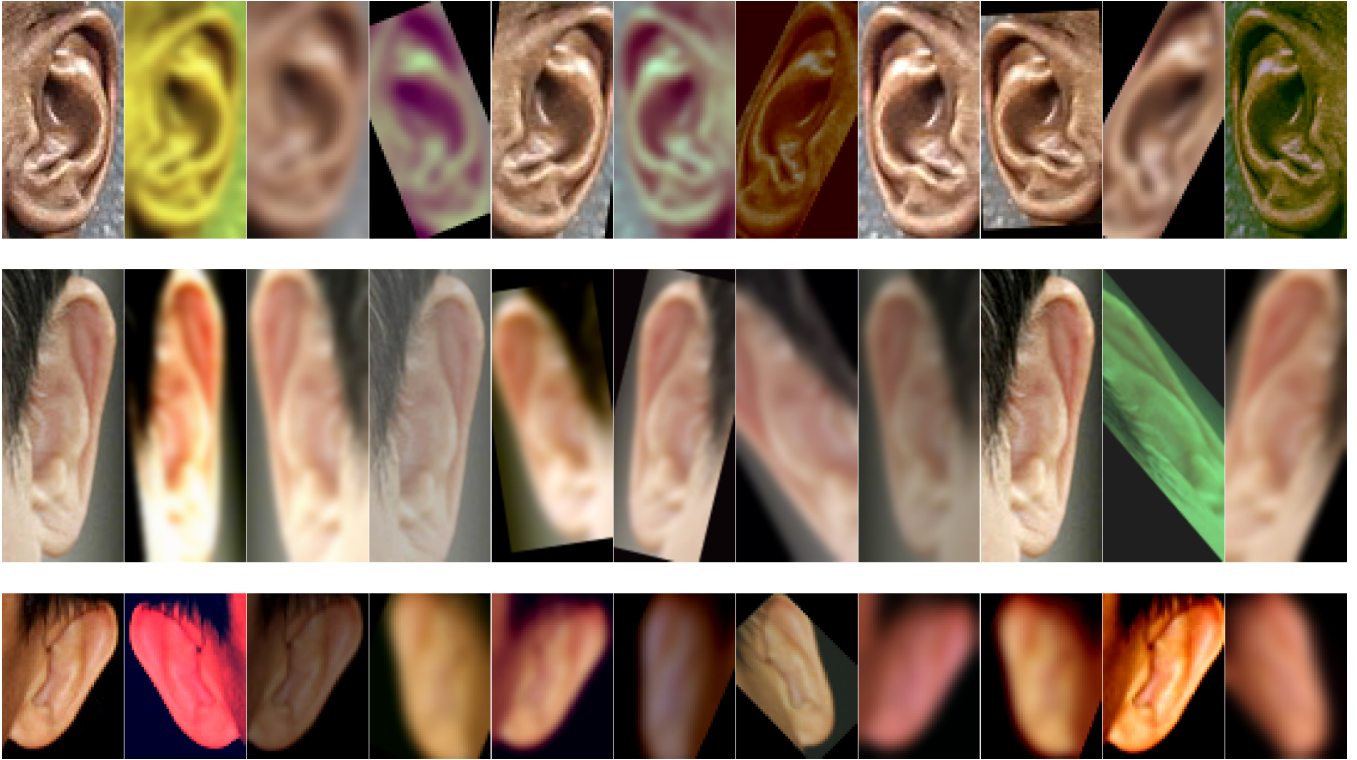


Fig. 1: Augmentation examples: the left most image in each row represents the original image from our dataset, while the following 10 images are the augmented variations. Note the extend of variability added to our training data by the augmented images.

exception here is the last fully-connected layer that needs to be learned from scratch and the softmax classifier at the top of the network that provides the output for our closed-set recognition problem. Below, we outline details of the selective training for each considered model architecture:

- AlexNet and VGG-16: For both architecture, we fix all models layers learned on the ImageNet data and initialize the first two fully connected layers with weights learned on ImageNet. Using the available training data for our problem, we fine tune these two layers and learn the third fully connected layer (which represents a proxy for the application-dependent softmax classifier) from scratch.
- SqueezeNet: Since this model has no fully-connected layers, we fine tune all layers except the last convolutional layer that has a different number of classes and therefore requires full learning. All model parameter are again initialized using the ImageNet data.

### B. Data augmentation

Both of our training strategies – full model learning as well as selective learning – require a sufficient amount of data for the training procedure [2]. In order to satisfy this requirement, we perform augmentations of the original dataset with translations, mild rotations, and color variations using the `Imgaug` tool (<http://github.com/aleju/imgaug>), where the data-transformations are performed (or not) with a

50% chance. Below is a list of the augmentation procedures we used to increase the amount of available training data.

- horizontal flipping,
- trimming 0–10% of images on each side,
- Gaussian blurring with  $\sigma$  0–3.0,
- addition of Gaussian noise with scale 0–0.2,
- brightness reduction/increase of pixel intensities by a value of 10 (over all color channels or over a single channel),
- contrast increase/decrease of up to 50% (over all color channels or over a single channel),
- rotation of up to  $45^\circ$  in both directions,
- scale increase/decrease of up to 20%.

Some sample augmentations are shown in Fig. 1. Here, the left most image in each row shows an example of an original image from our dataset, while the rest represents synthetic augmentations.

## IV. EXPERIMENTAL SETUP AND RESULTS

In our experiments we report the results for all strategies to CNN-model training outlined in the previous section. We present results for full-model and selective-model learning for all three architectures, but explore the impact of aggressive data augmentation only for the model with the least amount of parameters, i.e., SqueezeNet.

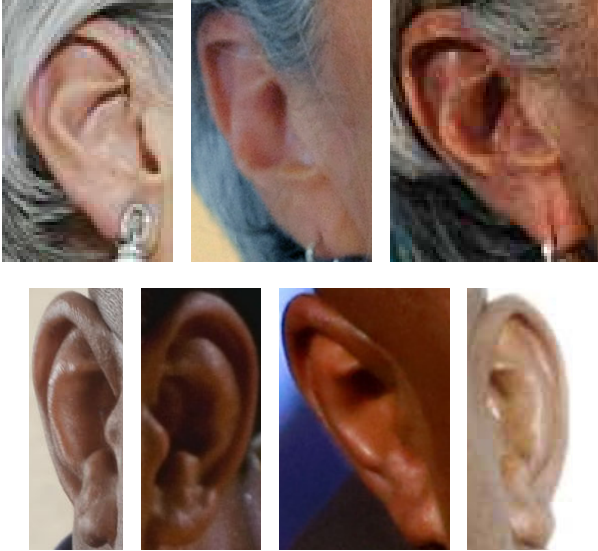


Fig. 2: Sample images from the dataset. Images in each row correspond to one subject from the dataset. Note the extent of variability present in the images.

### A. Dataset

For our experiments we compose a dataset of unconstrained ear images by merging the recently introduced AWED and CVLED datasets [19]. To have more data available to work with, we collect an additional 500 images of 50 subject from the Internet. The combined dataset contains 2304 images (1000 from AWED, 804 from CVLED, and 500 newly collected) of 166 persons.

All images are collected in a consistent manner from the internet using web crawlers and a subsequent manual inspection. Since the main purpose of the collected images is not ear-recognition research, the images contain realistic variations and present a challenging task to automatic ear recognition technology. Variations that can be found in the dataset are across gender, head rotations, race, presence of occlusions and alike. A few sample images from the dataset are shown in Figs. 1 in 2.

### B. Performance metrics & protocols

We perform identification experiments with a closed set experimental protocol. This means that for each image in the dataset the CNN-based recognition model needs to predict to which of the 166 classes the input image corresponds to. Based on that we report the following performance curves and metrics for the identification experiments:

- Cumulative Match-score Curves (CMC),
- Rank-1 and Rank-5 recognition rates,
- Area Under the CMC Curve (AUCMC).

For the experiments, the dataset is split into train and test sets in a ratio of 60% vs. 40%, respectively. Splitting is done for each subject separately: 60% of images of a given subject are randomly selected for the train set, the rest for the test set. This means that in the train set there are 1,383 images and

TABLE I: Performance metrics showing the effect of data augmentation on the recognition performance. Results here are shown for the SqueezeNet architecture only. Note how the increase in available training data improves performance.

#Iter.	#Augm.	Rank-1 [%]	Rank-5 [%]	AUCMC [%]
10k	0	36.37	57.33	89.91
	10	46.04	69.49	94.10
	100	56.68	77.09	96.09
20k	0	40.72	61.45	90.93
	10	56.57	74.92	93.88
	100	61.67	79.26	95.62
30k	0	41.26	61.45	90.96
	10	56.89	75.24	93.85
	100	61.89	80.46	95.48
40k	0	41.26	61.45	90.96
	10	56.89	75.46	93.86
	100	62.00	80.35	95.51
50k	0	41.26	61.45	90.96
	10	57.00	75.46	93.86
	100	62.00	80.35	95.51

921 are in the test set. All results are reported in identification experiments with the 921 test images.

### C. Results

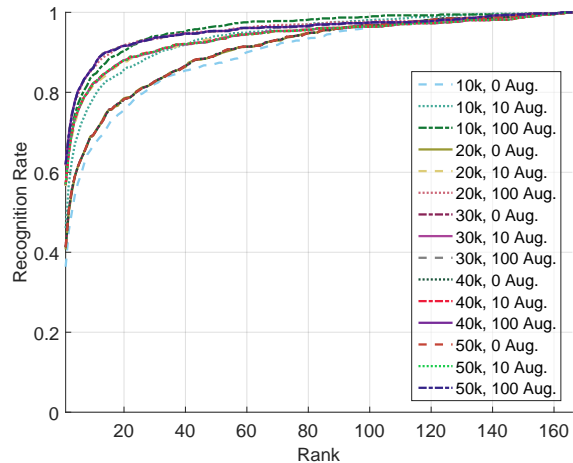
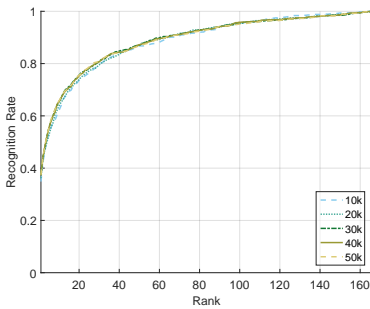


Fig. 3: CMC curves showing the effect of data augmentation on the recognition performance. Here we show only results for the SqueezeNet architecture. The figure legend provides information about the number of training iterations and augmentation factor used, i.e., (#Iter, factor Aug.).

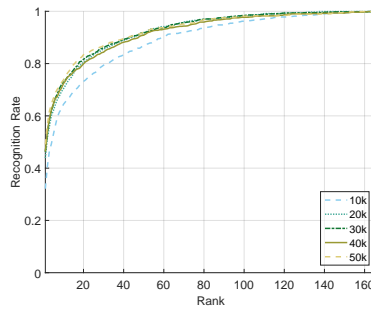
In the first series of experiments we evaluate the impact of aggressive data augmentation on CNN-based model training. For these experiments we consider only the SqueezeNet architecture, which is the fastest to train due to the relatively small number of model parameters. We perform experiments without data augmentation with the initial set of 1383 training images (i.e., an augmentation factor of 0), with 10-times the original training data (i.e., an augmentation factor of 10) and with 100-times the original training data (i.e., an augmentation factor of 100). We observe the recognition

TABLE II: Performance metrics for selective learning (bottom) and full model learning (top). The tables show results for different model architectures: (a) AlexNet, (b) VGG-16, and (c) SqueezeNet. The best results for each model architecture are presented in italic and the best overall results are marked bold.

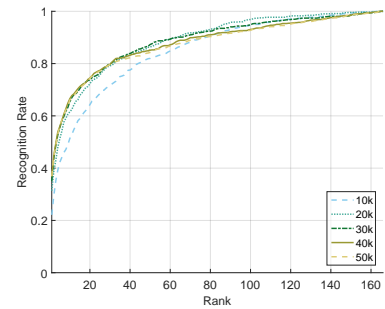
(a) AlexNet results				(b) VGG-16 results				(c) SqueezeNet results			
# Iter.	Rank-1 [%]	Rank-5 [%]	AUCMC [%]	# Iter.	Rank-1 [%]	Rank-5 [%]	AUCMC [%]	# Iter.	Rank-1 [%]	Rank-5 [%]	AUCMC [%]
Full model learning				Full model learning				Full model learning			
10k	34.85	53.31	88.74	10k	32.14	52.88	89.23	10k	22.15	42.24	85.43
20k	36.16	52.66	88.89	20k	43.87	62.00	92.63	20k	31.38	51.25	88.97
30k	37.57	55.48	89.19	30k	46.25	64.39	92.87	30k	35.07	55.48	88.87
40k	37.24	55.48	89.16	40k	46.36	64.60	92.16	40k	36.81	55.92	88.05
50k	37.46	55.37	89.08	50k	49.08	66.67	92.99	50k	36.92	56.03	87.58
Selective model learning				Selective model learning				Selective model learning			
10k	46.15	67.32	94.04	10k	48.10	67.97	94.14	10k	56.68	77.09	<b>96.09</b>
20k	49.29	69.60	94.48	20k	49.19	70.03	94.45	20k	61.67	79.26	95.62
30k	49.19	69.92	94.55	30k	50.27	70.90	94.66	30k	61.89	<b>80.46</b>	95.48
40k	49.51	69.71	94.54	40k	51.14	71.77	94.78	40k	<b>62.00</b>	80.35	95.51
50k	49.51	69.82	94.57	50k	51.25	71.99	94.81	50k	<b>62.00</b>	80.35	95.51



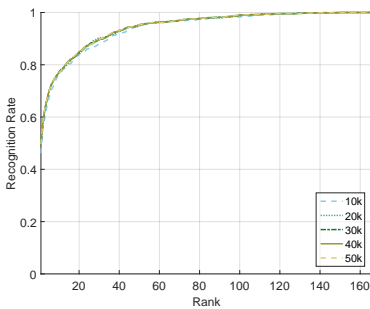
(a) AlexNet – full model learning



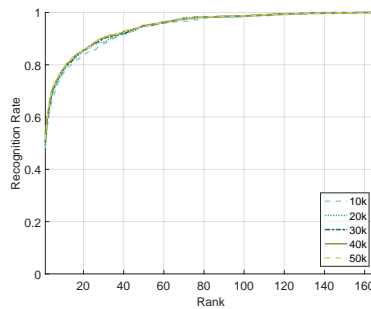
(b) VGG – full model learning



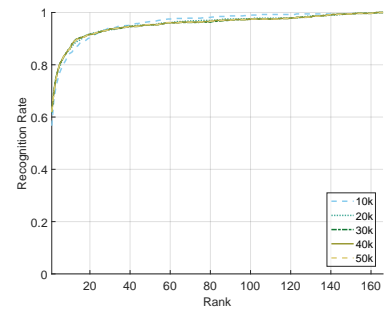
(c) SqueezeNet – full model learning



(d) AlexNet – selective model learning



(e) VGG – selective model learning



(f) SqueezeNet – selective model learning

Fig. 4: CMC curves generated based on different learning strategies. The first column shows curves for the AlexNet architecture, the second column shows results for the VGG architecture and the third column shows results for the SqueezeNet architecture. The upper row depicts results for full model learning and the lower for selective model learning. The best results are obtained with the SqueezeNet architecture with selective model learning. The results are best viewed in color.

performance of SqueezeNet trained with selective learning after 10,000, 20,000, 30,000, 40,000 and 50,000 training iterations. The results on the test set are presented in Table I and Fig. 3. After each group of iterations the augmentations prove to be very effective, increasing Rank-1 recognition rate by more than 20 percentage points when comparing no augmentation and the augmentation with a factor of 100. As expected, the highest augmentation factor results in the best overall performance. We, therefore, report all following results for experiments with an augmentation factor of 100.

In the next series of experiments, we explore our training

strategies on all three model architectures. We fix the data augmentation factor to 100 for all tests and vary only the number of training iterations from 20,000 to 50,000. The results for AlexNet in Table II(a) show that performance steadily increases as the number of iterations increases and starts stagnating after 30,000 iterations. However, while the trend is the same for both of our training strategies (full model learning and selective model learning) selective model learning significantly outperforms the full model learning approach with a 49.51% Rank-1 recognition rate vs. 37.46% for full model learning. The same conclusion holds for the

Rank-5 and AUCMC performance metrics with 69.82% vs. 55.37% and 94.57% vs. 89.08%, respectively. The two plots in the first column of Fig. 4 show CMC curves and again support the conclusion, where the CMC curves for selective model learning are higher for all iterations.

For VGG-16 the selective model learning again outperforms full model learning, as shown in Table II(b) and the two plots in the middle column of Fig. 4, but by a small margin of 2.17 percentage points for the Rank-1 recognition rate. The selective model learning on VGG-16 outperforms AlexNet with a Rank-1 recognition rate of 51.25% compared to the 49.51% of AlexNet. The best Rank-5 recognition rate and AUCMC values for the VGG-16 model are 71.99% and 94.81%, respectively, and are higher than with AlexNet.

SqueezeNet, same as AlexNet and VGG-16, performs better with selective model learning than full model learning, which again suggests that good parameter initialization is crucial for model training and helps reduce the need for large amounts of training data. The results for the SqueezeNet architecture are shown in Table II(c) and the CMC plots in the last column of Fig. 4. If we look at the results for full model learning only, SqueezeNet does not outperform VGG-16 with a Rank-1 recognition rate of 36.92% compared to 49.08% for VGG-16. However, it outperforms both, AlexNet and VGG-16, under the selective model learning strategy. The model results in a Rank-1 recognition rate of 62.00% vs. 49.51% and 51.25% for AlexNet and VGG-16, respectively.

Overall, the results suggest that selective model learning should be preferred to full model learning when only limited training data is available. While the results obtained with full-model learning were better than we initially expected, the training procedure proved much more difficult than with selective model learning. To find a good parameter configuration for the considered architectures with full model learning, we needed to use hyper-parameter optimization, test different parameter initialization schemes and repeat the training procedure several times. With selective model learning, on the other hand, the training procedure was straight forward and quickly converged to a good model configuration.

In the last series of identification experiments, we compare some existing state-of-the-art techniques to the best architectures obtained with the full-model-learning and selective-model-learning strategies. Here, we select the techniques to be included in our experiments based on recent publications evaluating the performance of various ear recognition techniques, i.e. [19], [20]. The state-of-the-art techniques included in our experiments are local techniques implemented in the AWE toolbox [19], which exploit the following descriptors: LBP [25], POEM [27], HOG [38], BSIF [39], LPQ [40], RILPQ [26], and DSIFT [41]. The results of this series of experiments are presented in Table III and Fig. 5. As can be seen, even the fully learned VGG-16 model significantly outperforms the best performing local approach, i.e., based on the Histogram-of-Oriented-Gradients (HOG) descriptor, with 49.08% vs. 34.64%, respectively. Furthermore, with selective learning SqueezeNet outperforms both

TABLE III: Performance metrics for the top performing CNN-based approaches for the full-model-learning and for the selective-model-learning strategies vs. some of the state-of-the-art feature extraction approaches. Results for the top performing local techniques are presented in *italic*, the overall best performance is marked **bold**.

# Iter.	Rank-1 [%]	Rank-5 [%]	AUCMC [%]
LBP [25]	28.12	45.17	76.40
POEM [27]	33.12	48.53	77.34
HOG [38]	<i>34.64</i>	<i>52.88</i>	<i>81.07</i>
BSIF [39]	32.25	47.23	77.69
LPQ [40]	29.21	44.95	76.49
RILPQ [26]	28.56	44.41	76.76
DSIFT [41]	27.69	42.67	75.74
Full learning (VGG-16)	49.08	66.67	92.99
Sel. learning (SqueezeNet)	<b>62.00</b>	<b>80.35</b>	<b>95.51</b>

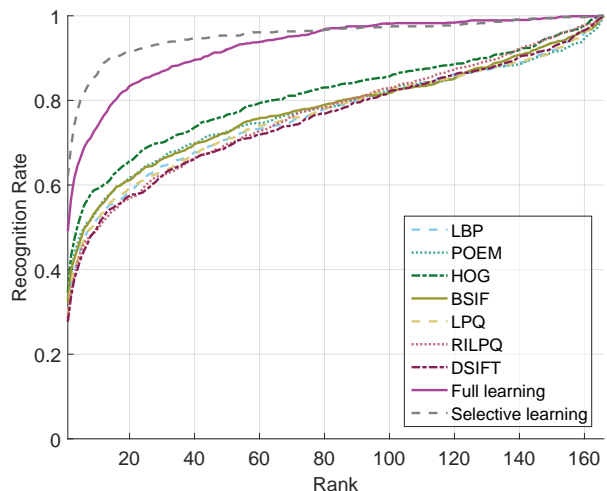


Fig. 5: CMC curves comparing state-of-the-art feature extraction methods to our best performing CNN-based approaches for full model learning (VGG-16) and for selective model learning (SqueezeNet).

HOG and POEM (which are the two best performing local techniques in our experiments) by a significant margin: 62.00% Rank-1 recognition rate compared to 34.64% for HOG and 33.12% for POEM. The same holds for Rank-5 recognition rates 92.99% and AUCMC 95.51% compared to 52.88% and 81.07% for HOG and 48.53% and 77.34% for POEM.

## V. CONCLUSION

In this work we studied the problem of training CNN-based models for (closed set) ear recognition using limited training data. We investigated different strategies towards model training and were able to build a model that improved on the best performing state-of-the-art technique (based on HOG descriptors) included in our comparison by close to 30% in terms of the Rank-1 recognition rate.

The best model we were able to build was based on the SqueezeNet architecture. The model was initialize with



parameters learned with the ImageNet data and then fine tuned using a limited set of 1383 ear images of 166 classes that were augmented by a factor of 100.

In terms of future work, our goal will be to develop CNN-based models for open-set recognition problems and move beyond the close-set protocol explored in this paper. The difficulty with this approach is that the softmax layer at the top of the networks will have to be removed and image descriptors computed at the fully connected layers will have to be used as image representations. This task is typically more challenging than closed set recognition, as the computed descriptors need to be representative of unseen image classes (or identities), which commonly induces the need for even more training data.

## REFERENCES

- [1] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," *arXiv preprint arXiv:1412.7062*, 2014.
- [2] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "Decaf: A deep convolutional activation feature for generic visual recognition," in *Icml*, vol. 32, 2014, pp. 647–655.
- [3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 779–788.
- [4] D. Erhan, C. Szegedy, A. Toshev, and D. Anguelov, "Scalable object detection using deep neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2147–2154.
- [5] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *European Conference on Computer Vision*. Springer, 2014, pp. 184–199.
- [6] Z. Cui, H. Chang, S. Shan, B. Zhong, and X. Chen, "Deep network cascade for image super-resolution," in *European Conference on Computer Vision*. Springer, 2014, pp. 49–64.
- [7] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2016.
- [8] J. Hu, J. Lu, and Y.-P. Tan, "Discriminative deep metric learning for face verification in the wild," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1875–1882.
- [9] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International Workshop on Similarity-Based Pattern Recognition*. Springer, 2015, pp. 84–92.
- [10] Y. Guo, Y. Liu, A. Oerlemans, S. Lao, S. Wu, and M. S. Lew, "Deep learning for visual understanding: A review," *Neurocomputing*, vol. 187, pp. 27–48, 2016.
- [11] J. Guo and S. Gould, "Deep cnn ensemble with data augmentation for object detection," *arXiv preprint arXiv:1506.07224*, 2015.
- [12] R. Dellana and K. Roy, "Data augmentation in cnn-based periocular authentication," in *Information Communication and Management (ICICM), International Conference on*. IEEE, 2016, pp. 141–145.
- [13] G. Hu, X. Peng, Y. Yang, T. Hospedales, and J. Verbeek, "Frankenstein: Learning deep face representations using small data," *arXiv preprint arXiv:1603.06470*, 2016.
- [14] K. Grm and V. Struc, "Deep face recognition for smart surveillance applications," *IEEE Intelligent Systems*, p. in press, 2017.
- [15] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [17] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [18] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size," *arXiv:1602.07360*, 2016.
- [19] Z. Emersic, V. Struc, and P. Peer, "Ear Recognition: More Than a Survey," *Neurocomputing*, 2017.
- [20] A. Pflug and C. Busch, "Ear biometrics: a survey of detection, feature extraction and recognition methods," *IET biometrics*, vol. 1, no. 2, pp. 114–129, 2012.
- [21] D. J. Hurley, M. S. Nixon, and J. N. Carter, "Automatic ear recognition by force field transformations," 2000.
- [22] L. Yuan, Z.-c. Mu, Y. Zhang, and K. Liu, "Ear recognition using improved non-negative matrix factorization," in *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, vol. 4. IEEE, 2006, pp. 501–504.
- [23] H.-J. Zhang, Z.-C. Mu, W. Qu, L.-M. Liu, and C.-Y. Zhang, "A novel approach for ear recognition based on ica and rbf network," in *Machine Learning and Cybernetics, 2005. Proceedings of 2005 International Conference on*, vol. 7. IEEE, 2005, pp. 4511–4515.
- [24] D. G. Lowe, "Object recognition from local scale-invariant features," in *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, vol. 2. Ieee, 1999, pp. 1150–1157.
- [25] Y. Guo and Z. Xu, "Ear recognition using a new local matching approach," in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*. IEEE, 2008, pp. 289–292.
- [26] V. Ojansivu, E. Rahtu, and J. Heikkilä, "Rotation invariant local phase quantization for blur insensitive texture analysis," in *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*. IEEE, 2008, pp. 1–4.
- [27] N.-S. Vu and A. Caplier, "Face recognition with patterns of oriented edge magnitudes," in *European conference on computer vision*. Springer, 2010, pp. 313–326.
- [28] L. Nanni and A. Lumini, "A multi-matcher for ear authentication," *Pattern Recognition Letters*, vol. 28, no. 16, pp. 2219–2226, 2007.
- [29] L. Jacob and G. Raju, "Ear recognition using texture features—a novel approach," in *Advances in Signal Processing and Intelligent Recognition Systems*. Springer, 2014, pp. 1–12.
- [30] A. Morales, M. Diaz, G. Llinas-Sanchez, and M. A. Ferrer, "Earprint recognition based on an ensemble of global and local features," in *Security Technology (ICCST), 2015 International Carnahan Conference on*. IEEE, 2015, pp. 253–258.
- [31] A. Abaza, A. Ross, C. Hebert, M. A. F. Harrison, and M. S. Nixon, "A survey on ear biometrics," *ACM computing surveys (CSUR)*, vol. 45, no. 2, p. 22, 2013.
- [32] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li, "Imagenet large scale visual recognition challenge," *CoRR*, vol. abs/1409.0575, 2014. [Online]. Available: <http://arxiv.org/abs/1409.0575>
- [33] S. Ren, K. He, R. B. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," *CoRR*, vol. abs/1506.01497, 2015. [Online]. Available: <http://arxiv.org/abs/1506.01497>
- [34] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *arXiv preprint arXiv:1511.00561*, 2015.
- [35] V. Badrinarayanan, A. Handa, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling," *arXiv preprint arXiv:1505.07293*, 2015.
- [36] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [37] P. L. Galdámez, W. Raveane, and A. G. Arrieta, "A brief review of the ear recognition process using deep neural networks," *Journal of Applied Logic*, 2016.
- [38] N. Damer and B. Führer, "Ear recognition using multi-scale histogram of oriented gradients," in *Intelligent Information Hiding and Multimedia Signal Processing (IHH-MSP), 2012 Eighth International Conference on*. IEEE, 2012, pp. 21–24.
- [39] J. Kannala and E. Rahtu, "Bsfif: Binarized statistical image features," in *Pattern Recognition (ICPR), 2012 21st International Conference on*. IEEE, 2012, pp. 1363–1366.
- [40] V. Ojansivu and J. Heikkilä, "Blur insensitive texture classification using local phase quantization," in *International conference on image and signal processing*. Springer, 2008, pp. 236–243.
- [41] K. Dewi and T. Yahagi, "Ear photo recognition using scale invariant keypoints," in *Computational Intelligence*, 2006, pp. 253–258.