

Received December 12, 2021, accepted January 13, 2022, date of publication January 25, 2022, date of current version February 2, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3146059

Prototype Memory for Large-Scale Face Representation Learning

EVGENY SMIRNOV¹, NIKITA GARAEV¹, VASILII GALYUK¹, AND EVGENY LUKYANETS²

¹Speech Technology Center, 194044 Saint Petersburg, Russia

²Department of Information Technologies and Programming, ITMO University, 197101 Saint Petersburg, Russia

Corresponding author: Evgeny Smirnov (smirnov-e@speechpro.com)

This work was supported by the Analytical Center for the Government of the Russian Federation (IGK 000000D730321P5Q0002) under Agreement 70-2021-00141.

ABSTRACT Face representation learning using datasets with a massive number of identities requires appropriate training methods. Softmax-based approach, currently the state-of-the-art in face recognition, in its usual “full softmax” form is not suitable for datasets with millions of persons. Several methods, based on the “sampled softmax” approach, were proposed to remove this limitation. These methods, however, have a set of disadvantages. One of them is a problem of “prototype obsolescence”: classifier weights (prototypes) of the rarely sampled classes receive too scarce gradients and become outdated and detached from the current encoder state, resulting in incorrect training signals. This problem is especially serious in ultra-large-scale datasets. In this paper, we propose a novel face representation learning model called Prototype Memory, which alleviates this problem and allows training on a dataset of any size. Prototype Memory consists of the limited-size memory module for storing recent class prototypes and employs a set of algorithms to update it in appropriate way. New class prototypes are generated on the fly using exemplar embeddings in the current mini-batch. These prototypes are enqueued to the memory and used in a role of classifier weights for softmax classification-based training. To prevent obsolescence and keep the memory in close connection with the encoder, prototypes are regularly refreshed, and oldest ones are dequeued and disposed of. Prototype Memory is computationally efficient and independent of dataset size. It can be used with various loss functions, hard example mining algorithms and encoder architectures. We prove the effectiveness of the proposed model by extensive experiments on popular face recognition benchmarks.

INDEX TERMS Deep neural networks, face recognition, representation learning, softmax acceleration.

I. INTRODUCTION

Face recognition is one of the most established technologies [1], [2] of computer vision. With the help of modern deep neural architectures [3] and large-scale training datasets [4], [5], current face recognition models [6]–[9] perform at super-human levels [10] on the million-scale face identification and verification tasks [11]. However, there are still a number of open problems in face recognition, such as large pose and age variations [12]–[15], racial bias [16], unconstrained face image conditions [17], usage of heavy makeup, disguise [18], and face masks [19].

The most straightforward way to combat these problems is a construction of large and (preferably) balanced training datasets with sufficient diversity of face image variations.

The associate editor coordinating the review of this manuscript and approving it for publication was Xinyu Du¹.

Indeed, the best performing models in popular face recognition benchmarks are trained on a datasets, containing hundreds of thousands [5], [8], [20] or even millions [21], [22] of persons.

Training on the datasets of a million-person scale provides a set of challenges. One of them is the choice of training method. First possible way to train a model on a dataset with a large number of persons is to rely on a triplet-based training scheme [21]. It is scalable, but usually slow and less accurate than alternatives like softmax-based methods [6], [7], [23].

Softmax-based methods consider face recognition as a classification task, where persons are classes, and the model is trained to classify an image as belonging to one of the classes in the training set. These methods are proven to perform well, but they have a problem with scaling. Computation and GPU memory requirements of these methods are dependent on a number of classes in the dataset. Small and

medium-sized [24]–[26] datasets with less than 100,000 classes could be used for softmax-based methods efficiently, but training on the datasets with millions of classes becomes too computationally demanding or even infeasible.

To overcome the difficulties of large dataset softmax-based training, several methods were proposed [5], [27]–[29]. These methods are based on the “sampled softmax” approach, where only a small subset of the total set of classes is used in the current training iteration. The advantage of this approach is that there is much less computation and GPU memory involved compared to the “full softmax” training. Only selected class prototypes (classifier weights) are moved to the GPU memory to perform training iteration. All other prototypes stay in RAM or in other large-capacity memory. Whereas the “full softmax” approach has the advantage of more complete information available at the loss computation stage, comparable results could be achieved with “sampled softmax” using smart selection of sampled classes [27], [29], [30] and large enough training datasets [5].

“Sampled softmax”-based methods are good solutions for some problems in face recognition model training on large-scale datasets, but they also have a set of disadvantages. First, they still need to store prototypes of all classes in memory. For multi-million-scale datasets and large prototype size it could become an unnecessary bottleneck. Also it means that the dataset must be fixed during the training process, and no new classes could be added. Second, operations of sampled prototype transfer from RAM to GPU memory and back add extra computational overhead. Third, class prototypes, which are not sampled in the current iteration, do not get gradients, and over time become detached from the encoder and begin to represent their class incorrectly in the embedding space. It leads to inaccurate training signals and subsequent performance degradation. We refer to this situation as a problem of “prototype obsolescence”. It is even more dangerous in large-scale datasets, where the frequency of sampling each individual class is very low.

In this paper, we propose Prototype Memory - a novel face representation learning model, suitable for training on a dataset, regardless of the number of identities in it. Prototype Memory is based on the idea of “online” class prototype generation. New class prototypes are approximated using groups of same-class exemplar embeddings in the current mini-batch, and stored in a memory of limited size. Prototypes in memory are used in a role of classifier weights to perform softmax-based training. Prototypes are updated using gradients and also refreshed with recent exemplar embeddings, when they are available. The memory is continuously filled with prototypes until it reaches its full capacity. In this case, to free the memory for new prototypes, the oldest ones are removed from the memory and disposed of.

Prototype Memory shares the advantages of other “sampled softmax”-based methods, but is free from many of their disadvantages. First, there is no need to store all class prototypes, only those in current Prototype Memory. This behavior allows us to use datasets with an unlimited number

of persons and train in class-incremental style. Second, the proposed approach has no overhead on RAM to GPU memory transfer operations: all Prototype Memory operations can be performed entirely on GPU. Third, class prototypes, used for training, are continuously updated and kept up-to-date by means of gradients and refreshing algorithm. Outdated prototypes are disposed of. This way we ensure that class prototypes are closely connected to the current encoder state and correctly represent their corresponding classes in the embedding space. It helps to solve the problem of “prototype obsolescence”.

Prototype Memory is orthogonal to loss functions, neural network architectures, hard example mining methods, etc. It can be combined with many different methods, including smart class [27], [29], [30] and example [29], [31] sampling algorithms, margin-enhanced softmax variants [7], [32], [33] and other state-of-the-art face recognition methods [34]–[38].

We summarize the contributions as follows:

- We propose Prototype Memory model for face representation learning, which can be used to efficiently train face recognition models on very large datasets.
- We prove the effectiveness of the proposed model with extensive experiments on different face recognition benchmarks.
- We present algorithms of hard example mining and knowledge distillation, suitable for Prototype Memory.

II. RELATED WORK

A. FACE RECOGNITION

Current state-of-the-art methods in the area of face recognition are based on the training of deep neural networks on large face image datasets [21], [39]. The usual pipeline of face recognition is to detect faces in image with a face detector [40], [41], align and crop them, and then pass through the encoder (deep neural network) to get face representations in a form of L_2 -normalized vectors (embeddings). Encoder is trained to represent different face images of the same person close to each other in the embedding space (according to the cosine similarity measure), and face images of two different persons - far from each other. There are two main ways to train an encoder to perform this task. First one is to group images in pairs [42] or triplets [21] and use their embeddings to compute gradients, pushing them to each other - for the same-class pairs, and apart from each other - for the pairs of different classes.

Second way to train proper face encoders is to present a task of face recognition in the form of softmax-based image classification [10] and apply additional constraints to the resulting image representations [43], [44] to ensure they possess necessary properties. Classifier weights in this approach could be viewed as class representatives (prototypes) in the embedding space, and exemplar embeddings are trained to move closer to their corresponding prototypes and away from the prototypes of other classes. With proper normalization, prototypes bear conceptual similarity to the

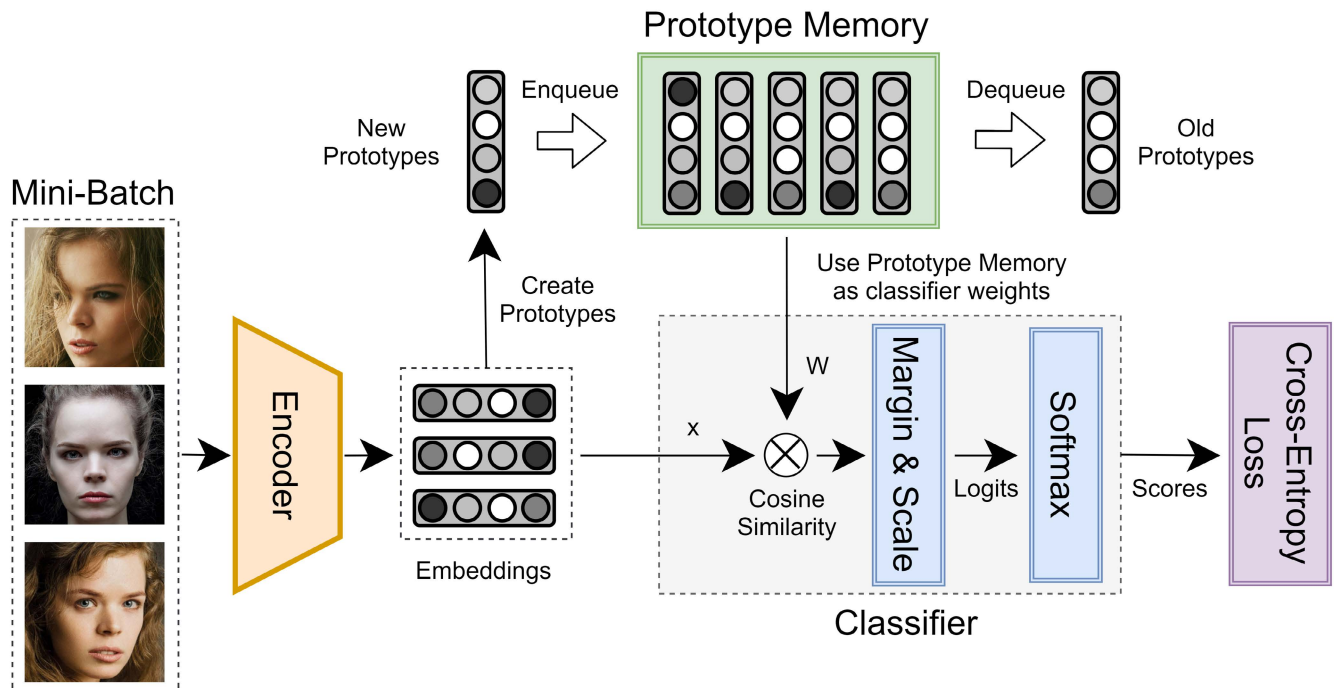


FIGURE 1. Scheme of the face representation learning with prototype memory. Mini-batch with several images per class is passed through the encoder, producing exemplar embeddings. Same-class embeddings are used to generate new class prototypes. These prototypes are enqueued to the memory module and used as the classifier weights in softmax classifier-based training. When the memory is full, the oldest prototypes are dequeued and disposed of.

class centers [45] in the embedding space. This approach is currently considered state-of-the-art and represented with such methods as SphereFace [23], ArcFace [7], CosFace [6], CurricularFace [32] and many others [29], [33], [46]–[49].

There are also many complementary methods proposed to build better face recognition models by promoting desired properties of the produced face representations, such as robustness to noisy labels [8] and low image resolution [50], invariance to age [51] and pose [52], ability to mitigate racial bias [53] and domain imbalance [35], [54], to improve the fairness of representations [55]. There are also methods, proposed to overcome the problems with the situations of difficult face appearance variations, like deliberately disguised faces [31], [56] or faces in medical masks [57].

All these methods are used to improve face recognition models, but the main key of success is the usage of large training datasets. For example, datasets, used to train state-of-the-art models in MegaFace challenge [11], are composed of hundreds of thousands [20] and even millions [22] of persons. To use this kind of datasets with the softmax-based approach there is a need of some kind of softmax acceleration methods.

B. SOFTMAX ACCELERATION

Utilization of softmax function is currently the main approach to perform classification tasks in the field of machine learning. In case of small number of classes, softmax is very fast, but when the number of classes reaches millions, softmax computation becomes a bottleneck. Recent works propose a number of ways to accelerate the computation of

softmax-based classification. One line of research is looking for the opportunities to accelerate softmax calculation with more efficient algorithms [58]–[60] and parallelization [5], [7], the other one is trying to approximate softmax [61], [62], split it into several parts [63], use softmax hierarchically [64], [65] or sample classes stochastically [66]–[69].

In the area of face recognition, the latter approach is the most promising. For example, a variant of “sampled softmax” approach, Selective Softmax [27] reduces the computation cost and memory demand of softmax computation by selecting only a subset of “active classes” for each mini-batch based on a set of dynamic class hierarchies, constructed on-the-fly. D-Softmax-K [28] randomly samples a subset of negative classes and uses two different loss functions for intra-class and inter-class training. Random Prototype Softmax and Dominant Prototype Softmax, proposed in [70], store a matrix of prototypes in the non-GPU memory and use different ways to select and move prototypes to the GPU memory to perform training in the bisample face recognition scenario. Hard Prototype Mining [29] is used to adaptively select a small number of hard prototypes, most similar to the samples in the mini-batch in each iteration. Only these prototypes are used in softmax calculation, thus accelerating the training.

Positive Plus Randomly Negative (PPRN) [5] is another softmax acceleration method. Its core idea is to keep a full classifier weight matrix in RAM and sample positive classes and a subset of random negative classes at each iteration to perform training. Provided with a large-scale training dataset,

this method of class sampling achieves the results, comparable to the “full softmax”-based training methods.

C. WEIGHT IMPRINTING

In the classic softmax-based training, classifier weights for all training classes are initialized with random values in the beginning of the learning process. However, with this kind of initialization all classes are need to be determined before the training, and their set cannot be altered after that. To achieve the possibility of inserting new classes to the trained network, [71] proposed a method of weight imprinting. With this method, embedding vectors of the examples of the new class can be used to initialize weights for this class, thereby extending the classifier. To generate a weight for a new class, embeddings of its examples are averaged, and the resulting vector is L_2 -normalized and inserted into the classifier. Another closely related model is Prototypical Networks [72]. They were proposed for the task of few-shot learning and use a support set of several examples per class to approximate the prototype of this class. Prototypes are calculated in the embedding space using the mean embedding of the support set. Classification is then performed by finding the nearest class prototype. There are a number of ways to use and adapt prototype-based learning in different areas [73]–[75]. In case of face recognition, class prototypes could be calculated using averaged and L_2 -normalized face image embeddings, belonging to the same class.

D. EXEMPLAR MEMORY

Exemplar memory is a way to improve the neural network training by keeping examples or their embeddings in memory between iterations and use them to get better training information. Exemplar memory was used in [76] to store the features of target domain and enforce invariance constraints for the task of person re-identification. Another proposed variant of exemplar memory is called Cross-Batch Memory (XBM) [77]. It was used to improve embedding learning by keeping embeddings from past iterations and using them to find better hard examples and learn better embeddings. Memory-based Jitter [78] enhances intra-class diversity for the tail classes in the long-tail learning problem by keeping in a memory bank different versions of exemplar features from previous iterations. BroadFace [34] is another variant of exemplar memory, proposed for face recognition. BroadFace stores a large number of exemplar embeddings from previous iterations and uses them to compute better training signals for the classifier weight matrix updates. Since the exemplar embeddings are computed with different versions of encoder from different training steps, they are suffering from the problem of obsolescence. To prevent it, authors proposed compensation function: using the difference between prototypes at current iteration and at the iteration, when the embedding was enqueued to the memory, calculate the compensation value and apply it to the embedding. Compensation function is helpful for embedding obsolescence, but it cannot be applied to prevent prototype obsolescence, because in contrast with

the BroadFace setting, in the “sampled softmax” situation there is no “actual” version of each class prototype at each iteration to get the compensation value from.

III. PROPOSED METHOD

In this section we propose Prototype Memory model for face representation learning. It combines the advantages of weight imprinting (online generation of class prototypes using groups of class exemplars), exemplar memory (keeping useful information between training iterations) and accelerated softmax methods (faster and more memory-efficient than “full softmax”) and could be used to train face recognition architectures on datasets of any size without the problem of prototype obsolescence.

A. MOTIVATION

Current state-of-the-art face recognition models [7] use softmax-based training strategy. When the dataset is very large (millions of persons), this type of strategy becomes too memory and computationally expensive: for the training with “full softmax”, all classifier weights should be placed in the GPU memory, and for each training example all classes in the classifier are needed to compute the loss function and gradients.

Current “sampled softmax”-based methods [5], [27]–[29], while solving some of the problems of “full softmax”, still have the need to keep all classifier weights somewhere, and suffer from the problem of “prototype obsolescence” - situations when the class prototype has not received gradients for so long time, that it does not represent its class correctly anymore. This problem is unavoidable in current methods if just a subset of classifier weights is updated at each iteration.

One more problem we want to solve is the inability of the current face recognition methods to use training datasets in class-incremental fashion, when the dataset could be updated with new classes at any time, without training problems.

B. PROTOTYPE MEMORY

The scheme of face representation learning with the proposed Prototype Memory model is presented in Fig. 1. Prototype Memory includes a memory module and algorithms for prototype generation, prototype refreshing and prototype disposal.

At each training iteration, several images of some class are passed through the encoder network, producing L_2 -normalized embeddings. These embeddings are used to generate a new prototype for the corresponding class. New class prototypes are enqueued to the Prototype Memory module. Prototypes in this module are later used as the classifier weights to perform usual softmax-based face recognition model training with some appropriate loss function.

Prototypes in the Prototype Memory are updated with the gradients. If the class, currently already placed in the Prototype Memory, gets new examples in the current mini-batch, it also uses their embeddings to “refresh” its prototype and move it to the start of the queue. When the Prototype Memory

is filled to its capacity, oldest prototypes, which haven't been refreshed with exemplar embeddings for many iterations and became outdated, are dequeued and disposed of to free the memory slots for the new ones.

Prototype Memory model has a set of hyperparameters:

- **Prototype Size (D)** - it is the size of the prototype vectors. This size must be the same as the size of the embedding vectors, produced by the encoder.
- **Memory Size (M)** - is the number of prototypes, which could be placed in the memory module.
- **Number of images per class (k)**, used for the prototype generation.
- **Refresh Ratio (r)** - this hyperparameter is responsible for the strength of the prototype refreshing.

More detailed description of the Prototype Memory algorithms is presented below and on Fig. 3.

1) PROTOTYPE GENERATION

Prototype generation scheme is presented at Fig. 3a. Mini-batch, containing the images of several classes, with the fixed number of images per class (set by a hyperparameter k), is passed through the encoder network. Resulting embeddings for each person (identity) are averaged, and the average embedding is L_2 -normalized:

$$\mathbf{P}_{new} = F_{norm}\left(\frac{\sum_{j=1}^k \mathbf{x}_j}{k}\right), \quad (1)$$

\mathbf{x}_j is the j -th exemplar embedding, produced by encoder, k is the number of images, used for prototype generation, and F_{norm} is L_2 -normalization function. The resulting vector \mathbf{P}_{new} is the new class prototype. New prototypes are calculated for each class in the mini-batch, and then enqueued to the Prototype Memory (Fig. 3b).

2) PROTOTYPE UPDATE AND REFRESHING

Prototype Memory module has fixed size, set by a hyperparameter M . New prototypes are enqueued to the Prototype Memory (and later dequeued from it) in order of appearance in training. When the prototypes are in Prototype Memory, they are used in a role of the classifier weights for usual softmax classifier-based training with a margin-based face-recognition-specific loss function like ArcFace [7]. These prototypes receive gradients and are updated according to them.

When the class has new examples in the current mini-batch, and already has a prototype in the Prototype Memory, additionally to the usual gradient-based update, it is also "refreshed" using current mini-batch embeddings with:

$$\mathbf{P}_{upd} = F_{norm}(r * \mathbf{P}_{new} + (1 - r) * \mathbf{P}_{mem}), \quad (2)$$

where \mathbf{P}_{upd} is updated prototype, \mathbf{P}_{mem} is a prototype from memory, \mathbf{P}_{new} is a prototype, calculated with the embeddings in the current mini-batch, F_{norm} means L_2 -normalization, and r is a hyperparameter (refresh ratio). Prototype refreshing is performed to force the prototypes in Prototype Memory to be

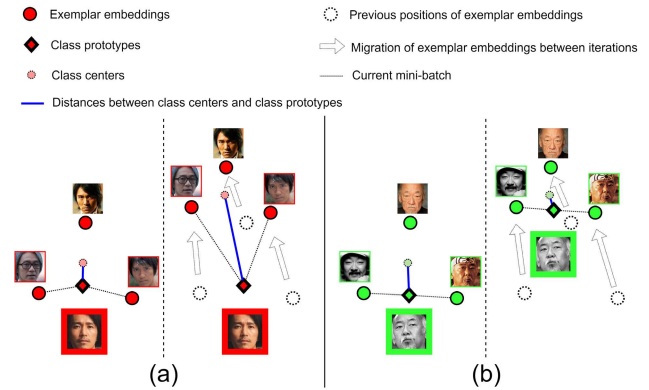


FIGURE 2. (a) Illustration of the prototype obsolescence problem in the "sampled softmax"-based training methods, (b) Illustration of the absence of prototype obsolescence problem in prototype memory.

up-to-date to the current encoder state. When the prototype is refreshed, it is moved to the start of the queue as if it is a new prototype.

3) PROTOTYPE DISPOSAL

Prototype disposal scheme is presented at Fig. 3b. When the Prototype Memory is full and has no free slots for new prototypes, there is a need to dispose of the oldest prototypes. These are the prototypes, which resided in the memory for the longest time, received no refreshing from their examples for a while and became outdated. Most likely, the "real" prototypes of these classes, if computed using the current encoder state and the images of this class, would be far from these prototypes in memory. To free the memory space for the new prototypes, these prototypes are dequeued and disposed of. Freed memory slots are then used for new prototypes.

C. SOLUTION TO PROTOTYPE OBSOLESCENCE

Prototype obsolescence problem of "sampled softmax"-based training methods is illustrated in Fig. 2a. Left part of the figure demonstrates a situation, when the class prototype is up-to-date, and the distance between prototype and class center (mean embedding for all class examples) is small. Right part of the figure shows a situation after a number of iterations, during which this class has not been sampled, and the prototype was not updated. Encoder has evolved, and now it produces exemplar embeddings in different points of the hypersphere. This resulted in a new position of the class center. Class prototype, however, stayed in the same place, and now the distance between prototype and class center became large. This prototype became obsolete, it is not representing the class correctly in the embedding space anymore.

Fig. 2b illustrates a similar situation, but for Prototype Memory. Instead of using the same old prototype from previous iterations, a new prototype is generated using exemplar embeddings in the current mini-batch. In this way it is kept up-to-date, even after a large number of iterations between the mini-batches, when the same class is sampled. Since new

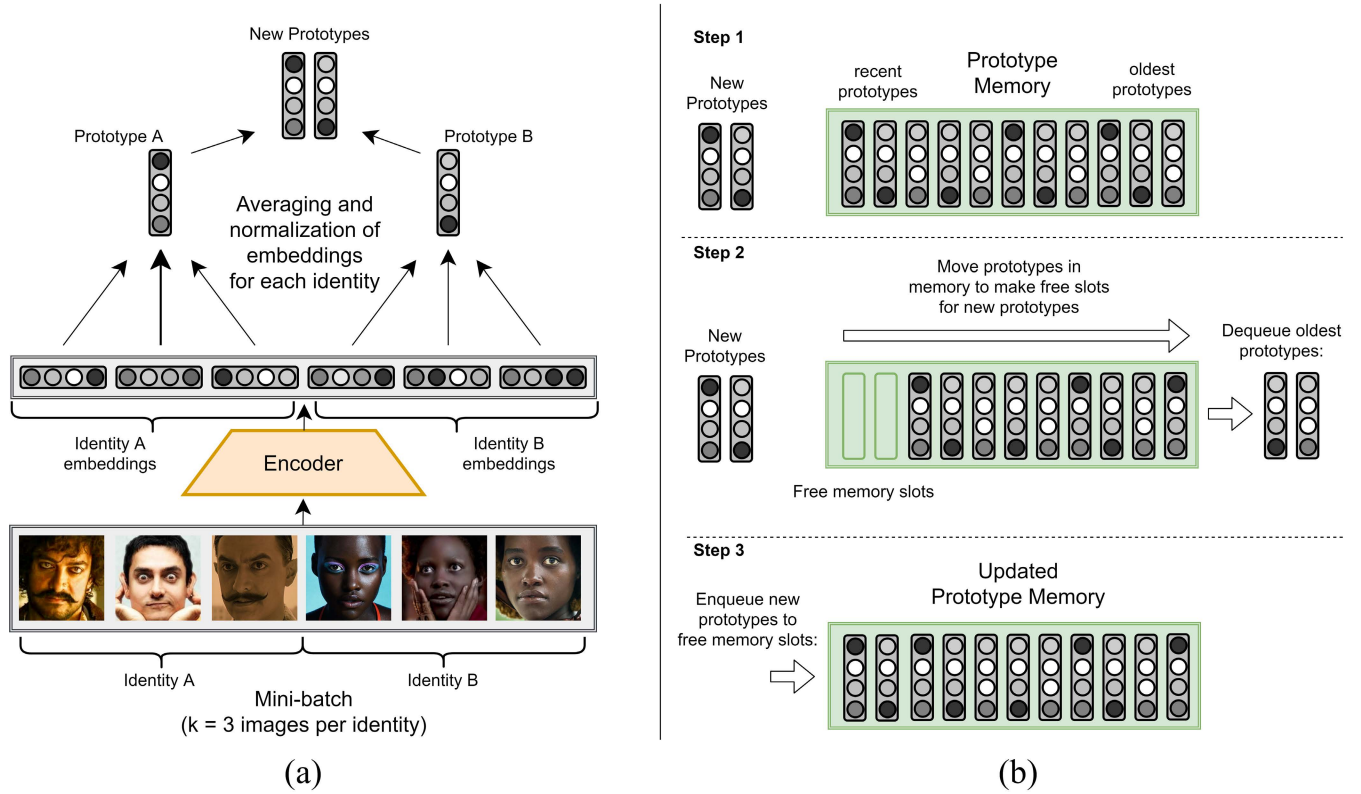


FIGURE 3. a) Scheme of prototype generation, b) Scheme of Prototype Memory update with new prototypes and the disposal of old prototypes.

class prototypes are generated using up-to-date encoder, the distance between class centers and prototypes stays small, so the problem of prototype obsolescence is solved.

D. TRAINING SPECIFICS

1) GROUP-BASED MINI-BATCH SAMPLING

For the prototype generation stage, to create prototypes closer to the corresponding class centers, we need to ensure that each class in the mini-batch has some minimum number of images k , sampled together. So, we need to sample images into the mini-batches accordingly. In this paper we utilize two different types of methods, used to sample mini-batches with groups of at least k images for each class:

- **Group-based iterate-and-shuffle:** First, images in the training dataset are randomly combined into the single-class groups of the same size k . Then these groups are shuffled and iteratively sampled to the mini-batch. When all groups are used, the process is repeated. With this kind of mini-batch sampling, we ensure that all images in the dataset will be sampled approximately equal number of times.
- **Group-based classes-then-images:** First, a number of classes are sampled from the training dataset. Then for each class a group of k random images sampled to the current mini-batch. With this kind of mini-batch sampling, we ensure that all classes in the dataset will be sampled approximately equal number of times.

These methods could be combined with algorithms of hard class mining [30] and hard example mining [31], or used together as parts of a composite mini-batch [56].

2) LOSS FUNCTIONS

Prototype Memory is a model, independent of the choice of the loss function [79]. Most loss functions, suitable for usual softmax-based face representation learning, could be also used to perform training with Prototype Memory. In our experiments, we use Large Margin Cosine Loss [6]:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\cos \theta_{y_i} - m)}}{e^{s(\cos \theta_{y_i} - m)} + \sum_{j=1, j \neq y_i}^M e^{s \cos \theta_j}}, \quad (3)$$

where N is a size of the mini-batch, M is the size of the memory module, y_i is the ground-truth class, θ_j is an angle between i -th exemplar embedding and j -th class prototype, s is a scale parameter, and m is margin.

We also performed experiments with D-Softmax [28] loss:

$$L_D = \frac{1}{N} \sum_{i=1}^N \log(1 + \frac{\epsilon}{e^{s \cos \theta_{y_i}}}) + \log(1 + \sum_{j \neq y_i}^M e^{s \cos \theta_j}), \quad (4)$$

where $\epsilon = e^{ds}$, d is a hyperparameter (optimization termination point), s is a scale parameter, N is a size of the mini-batch, M is the size of the memory module, y_i is the ground-truth class, and θ_j is an angle between i -th exemplar embedding

and j -th class prototype. Other similar loss functions [7], [32], [47] are applicable too.

E. HARD EXAMPLE MINING FOR PROTOTYPE MEMORY

Prototype Memory could be used together with hard example mining [31] and hard class mining methods like Doppelganger Mining [30], DP-Softmax [70] or HPM [29]. Careful selection of hard classes and examples is used to decrease the number of required training iterations and to achieve better performance in difficult face recognition scenarios. Most methods could be used for Prototype Memory without modifications, but some of them need to be adapted to the “sampled softmax” scenario.

1) MULTI-DOPPELGANGER MINING

Doppelganger Mining [30] is a hard class mining method, proposed for the “full softmax” classifier-based face representation learning scenario. Its main idea is to maintain a list with the most similar persons for each person in the training set. These most similar persons are called “doppelgangers”, and each person in the dataset has exactly one doppelganger in the list. Doppelgangers are updated at each training iteration, using softmax classification scores of non-target classes: class with the largest score becomes a new doppelganger for a class, which exemplar is classified. At the mini-batch generation stage, a doppelganger list is used to sample pairs of similar classes together.

Doppelganger Mining has demonstrated the ability to improve face recognition models, but it is based on the assumption that all training classes are simultaneously presented in the softmax classifier at each training iteration. It is true for “full softmax”-based training, but false for models, based on “sampled softmax”. In the latter case, only a subset of classes is used in the classifier at each iteration, so there is no guarantee that the most similar class (“global” doppelganger) is presented in the current classifier. As the model is forced to update the doppelganger list at each iteration, it may rewrite “global” doppelganger class in the list with a “local” one, by searching only in a subset of currently sampled classes. When the total number of classes in the dataset is large, and the number of sampled classes is small, there is a high chance, that doppelganger list will be filled with mostly “local” doppelgangers, missing the “global” ones, which are more preferable.

To prevent it, we propose Multi-Doppelganger Mining - a modification of Doppelganger Mining, adapted for the usage with “sampled softmax” models like Prototype Memory. With Multi-Doppelganger Mining, each class in the training dataset has a set of doppelgangers in the doppelganger list (instead of just one doppelganger per class as in the original version). Doppelgangers are added to this set using classifier scores in “sampled softmax”: non-target class with the largest classification score is added as a new doppelganger to the set. If a class in the doppelganger set is presented in the current “sampled softmax” classifier, but does not get the largest classification score, - it is removed from the

doppelganger set. At the stage of mini-batch generation, the doppelganger class is selected randomly from the doppelganger set.

2) HARDNESS-AWARE EXAMPLE MINING

In the “group-based classes-then-images” mini-batch generation strategy, when the classes are selected (randomly or using hard class mining algorithm), there is a need to select examples of these classes to be used in the mini-batch. While random sampling could be performed successfully in most cases, sometimes the number of examples per class is very large, and only a small portion of them possess enough utility for the training process.

To find the most useful examples and put them in the mini-batches more frequently, we propose a hardness-aware example mining method. Its main idea is to use cosine similarity scores, which are calculated between exemplar embeddings and class prototypes in the process of softmax-based classification, to measure the “hardness” of each example in the dataset, and then use these hardness values in the mini-batch generation stage to sample hard examples more frequently. Calculated hardness values are kept between training iterations and rewritten every time, when the example is classified again. Hardness values are calculated using:

$$H_i = 1 - \cos \theta_{y_i}, \quad (5)$$

where y_i is the ground-truth class for exemplar i , and θ_{y_i} is an angle between exemplar embedding and class prototype. Hardness values are initialized with $H_i = 2.0$. To control the amount of hard examples in the mini-batch, hyperparameter h (hardness ratio) is used. It determines the proportion of mini-batch examples, which are sampled using hardness values instead of random sampling. When $h = 0.0$, it means that all examples are sampled randomly, $h = 1.0$ means all examples are sampled using their hardness values, $h = 0.5$ means half of the mini-batch is sampled using hardness, and half of the mini-batch - randomly.

F. PROTOTYPE MEMORY KNOWLEDGE DISTILLATION

Knowledge distillation [80], [81] is a useful method of transferring knowledge from large teacher models to the smaller student models. There are several knowledge distillation methods, suitable for face recognition models [82]–[84], but they are not adapted to the Prototype Memory approach. For example, to initialize the training, [83] needs class centers for each class, calculated by teacher network. They could either be taken from the classifier weight matrix, or computed by averaging all teacher embeddings of each class in the dataset. This method requires memory space for prototypes of all dataset classes, thus limiting the ability of Prototype Memory to scale to the unlimited dataset sizes.

Here we propose Prototype Memory Knowledge Distillation (PMKD) - an approach of Prototype Memory-suitable knowledge distillation. Its essence is to calculate embeddings of the training dataset images with a teacher network (it could be done offline or online, in parallel with the training of the

student, on a dataset of any size), and use these embeddings instead of the produced by the student encoder to generate prototypes for the Prototype Memory.

IV. EXPERIMENTS

In this section we perform experiments with different Prototype Memory hyperparameters. We demonstrate the ability of Prototype Memory to solve the problem of prototype obsolescence. We also compare the performance of the proposed model with different “sampled softmax”-based alternatives and evaluate the effectiveness of the proposed hard example mining and knowledge distillation methods. Finally, we provide the comparison of the proposed Prototype Memory model with the state-of-the-art results on popular face recognition benchmarks.

A. IMPLEMENTATION DETAILS

1) DATASETS

For training we used Glint360k [5] dataset, containing 360,232 persons and 17M images, its cleaned variant Glint360k-M, containing 356,892 persons (we found identity duplicates and merged them) and 17M images, and also Glint360k-R with 353,658 persons and 16.3M images, where we removed identities, overlapping with testing datasets to perform more strict and fair evaluations [85]. For the experiments with prototype obsolescence prevention, we also used cleaned version of MS-Celeb-1M [4] dataset, containing 90,382 persons and 5.2M images. For testing we used datasets: LFW [86], CFP-FP [12], AgeDB-30 [14], CALFW [15], CPLFW [13], Trillion-Pairs [87] and MegaFace, original [11] and refined (R) version [7].

2) TRAINING SETTINGS

We use face images of size 112×112 , detected and cropped by RetinaFace detector [41], and employ ResNet-34 and ResNet-100 [88] architectures for the encoder. Final L_2 -normalized embeddings are of size $D = 256$ for ResNet-34 and $D = 512$ for ResNet-100 models. For experiments with Prototype Memory hyperparameters we used CosFace [6] loss with $s = 64$ and $m = 0.4$. The learning rate started from 0.1 and divided by 10 at 100k, 200k, 250k, 275k iterations finishing at 300k iterations. For the experiments with larger networks, we used PM-100: ResNet-100, pre-trained using Prototype Memory ($M = 200,000$, $r = 0.2$) and CosFace with $s = 64$ and $m = 0.4$. The learning rate started from 0.1 and divided by 10 at 200k, 400k, and 500k iterations, finishing at 540k. Mini-batch size is 512 for all models, mini-batch sampling is group-based iterate-and-shuffle.

3) TESTING SETTINGS

For testing we used L_2 -normalized average embedding of image and its horizontally flipped copy. For testing with the ResNet-100 model on the MegaFace dataset we used

L_2 -normalized concatenated embeddings of image and its horizontally flipped copy. For the metrics we employ verification accuracy on LFW, CFP-FP, AgeDB-30, CALFW and CPLFW datasets. We also perform evaluations on TrillionPairs (identification TPR@FAR = $1e-3$ and verification TPR@FAR = $1e-9$) and on MegaFace, original and refined (identification rank-1 and verification TPR@FAR = $1e-6$).

B. HYPERPARAMETER EFFECTS

We have performed the experiments with different hyperparameters of Prototype Memory.

1) NUMBER OF IMAGES PER CLASS (K) FOR PROTOTYPE GENERATION

We have performed experiments with the number of images per class (k) in the mini-batch. We used a ResNet-34 model with memory size of $M = 36,000$ and refresh ratio $r = 0.2$, trained on Glint360k with a mini-batch of size 512, and performed the testing on LFW, CFP-FP, AgeDB and Trillion-Pairs. The results are in Table 1.

TABLE 1. Effects of the number of images per class for prototype generation.

k	LFW	CFP-FP	AgeDB	Trillion-Pairs	
				Id	Ver
3	99.70	98.50	97.95	72.18	73.07
4	99.73	98.40	98.20	76.03	75.27
5	99.67	98.36	97.72	75.22	74.92

The best results were achieved with $k = 4$ images per class. Smaller value of k results in less accurately generated prototypes, larger k results in smaller number of different identities in each mini-batch. With larger mini-batch sizes and with datasets with many images per class it is reasonable to try larger values of k .

2) REFRESH RATIO

We have performed experiments with different Prototype Memory refresh ratios (r). $r = 0.0$ means prototypes in memory are not refreshed, $r = 1.0$ - prototypes in memory are completely replaced with new prototypes, $r = 0.2$ and $r = 0.5$ are other values, providing different levels of the refreshing strength. We used a ResNet-34 model with memory size of $M = 36,000$ and $k = 4$ images per class in the mini-batch, trained on Glint360k, and performed the testing on LFW, CFP-FP, AgeDB and Trillion-Pairs. The results are in Table 2.

TABLE 2. Effects of the refresh ratio (r).

r	LFW	CFP-FP	AgeDB	Trillion-Pairs	
				Id	Ver
0.0	99.73	98.47	97.70	73.76	72.95
0.2	99.73	98.40	98.20	76.03	75.27
0.5	99.73	98.36	97.37	66.15	65.80
1.0	99.45	95.27	95.80	53.98	53.41

The best results were achieved with $r = 0.2$, demonstrating the effectiveness of prototype refreshing. With $r = 0.0$ prototypes are getting updates from their class exemplars

only in a form of infrequent gradients, leading to a mild form of “prototype obsolescence”. Large r values refresh prototypes too strongly, resulting in the loss of valuable information, accumulated during training by the prototypes in memory. $r = 0.2$ is a value, achieving the balance between these two situations.

3) MEMORY SIZE

We have performed experiments with different memory sizes M for Prototype Memory. We used a ResNet-34 model with memory size ranging from 36,000 to 288,000, with other hyperparameters set as $r = 0.2$, and $k = 4$ images per class in the mini-batch of size 512. We have trained the models on Glint360k-M, and performed the testing on MegaFace (R) identification benchmark. The results are presented in Fig. 4.




FIGURE 4. Effect of the memory size M on MegaFace (R) rank-1 identification accuracy.

The best results were achieved with memory size of 216,000. With the increase of the memory size from 36,000 to 216,000, performance of the model improved. Presence of the sufficient number of prototypes in memory is essential for the construction of informative training signals. Large memory with correctly calculated prototypes provides a comprehensive model of the class distribution in the embedding space, which is used for accurate calculation of gradients to train precise face recognition models. However, it is hard to maintain the correctness of all prototypes in large-scale memory, when the encoder, which produces face representations, is evolving, and prototypes in the memory are updated with delays and imperfect information.

The decrease of the face recognition accuracy for larger memory sizes in our experiments illustrates the harmful impact of the infrequency of the target class exemplar-based gradient updates. With the large memory sizes, most gradient updates, received by the prototypes in memory, come from the non-target class examples, pushing these prototypes away from them in the embedding space. Target class exemplar-based updates, on the contrary, are rare and thus provide insufficient training signals to the prototypes to keep them in close connection with exemplar embeddings. Prototype

refreshing is helpful for this situation, but it is also dependent on the presence of target class examples in the mini-batch, and for large memory sizes its effect is limited. Therefore, to achieve good results, memory size should be large enough, but at the same time limited to a certain size.

C. EXPERIMENTS WITH PROTOTYPE OBSOLESCENCE

To measure the effect of prototype obsolescence, we calculated average distances between class prototypes and class centers for 100 longest unsampled classes for Positive Plus Random Negative (PPRN) model [5] and Prototype Memory. Class prototypes were taken from the weight matrix in RAM for the PPRN model, and generated using $k = 4$ random images of the considered class for Prototype Memory. Class centers were calculated as mean embeddings for the corresponding classes, using all class images, both for PPRN and Prototype Memory. We have performed experiments on the cleaned MS-Celeb-1M dataset with 90,382 classes and on the Glint-360k-R dataset with 353,658 classes. We used a softmax size of 36,000 (≈ 0.1 sampling rate) for both PPRN and Prototype Memory, and trained for 50,000 iterations. We used the ResNet-34 model as the encoder and CosFace with $s = 64$ and $m = 0.4$ for the loss function. Prototype Memory used $r = 0.2$ and $k = 4$. Results of the experiments are presented on Fig. 5.




FIGURE 5. Average cosine distances between class prototypes and class centers for 100 longest unsampled classes a) on MS-Celeb-1M dataset, without prototype obsolescence problem, b) on Glint-360k-R dataset, with prototype obsolescence problem.

On MS-Celeb-1M dataset (Fig. 5a), where the number of classes is moderate, prototype obsolescence is not present. After a few iterations, class prototypes for the longest unsampled classes move close to the class centers, both for PPRN and Prototype Memory. For the Glint-360k-R dataset (Fig. 5b), however, even after a large number of iterations the distance between class prototypes and class centers for PPRN remains considerably large. This discrepancy between prototypes and class centers is a variant of the problem of prototype obsolescence. It resulted in incorrect training signals, directing the updates in wrong directions. For the datasets of larger size this effect will be even more severe.

Unlike PPRN, the distance between class centers and prototypes for Prototype Memory remains insignificant both for MS-Celeb-1M and Glint-360k-R datasets. Thus we demonstrate that prototype obsolescence is not present in Prototype Memory. With the help of prototype refreshing algorithm, and due to the disposal and subsequent re-generation of the oldest

prototypes, and by the means of the limited size of memory module, prototypes in memory remain up-to-date and close to class centers at each training iteration, resulting in more accurate training signals and ability to train face recognition models regardless of the dataset size.

D. COMPARISON WITH OTHER “SAMPLED SOFTMAX”-BASED METHODS

We have performed the comparison with other “sampled softmax”-based methods - Positive Plus Random Negative (PPRN) [5] and D-Softmax-K [28]. We used a ResNet-34 model with memory size of 36,000 (≈ 0.1 sampling rate), $r = 0.2$ and $k = 4$ images per class. For PPRN we used a sampling rate of 0.1 and CosFace with $s = 64$ and $m = 0.4$. For D-Softmax-K we used $d = 0.9$, $s = 64$ and sampling rate of 0.1. Training was performed on Glint360k-M, testing on MegaFace, original and cleaned variants (the latter is more reliable). Results of the experiments are presented in Table 3.

TABLE 3. Comparison with other “sampled softmax”-based methods.

Method	MegaFace		MegaFace (R)	
	Id	Ver	Id	Ver
PPRN [5]	80.43	96.19	97.33	97.54
D-Softmax-K [28]	80.94	96.25	97.53	97.60
Prototype Memory	81.48	96.02	97.60	98.25

Prototype Memory outperformed other sampling-based methods in terms of face recognition accuracy. These results prove the effectiveness of the proposed model for learning face representations on large-scale training datasets.

Besides being accurate, Prototype Memory is also more memory-efficient than other softmax-based methods. In Table 4 we compare GPU and Non-GPU memory requirements of Prototype Memory, “full softmax”-based training, PPRN and D-Softmax-K. For the case of “full softmax”, full weight matrix of size $D \times N$ should be placed in GPU memory, where D is embedding size and N is total number of classes in the training dataset. There is no need for additional non-GPU memory. For PPRN and D-Softmax-K, only a small part of the full weight matrix is placed in GPU memory, however there is a need to keep the full weight matrix in non-GPU memory. The size of required GPU memory is $D \times M$, where $M < N$ is the number of classes in sampled softmax. The size of required non-GPU memory is $D \times N$.

For the case of Prototype Memory, there is no need to keep a full weight matrix in non-GPU memory, only $D \times M$ of GPU memory is required. Prototype Memory provides a memory-efficient way of training face representations. Memory requirements of the proposed model are independent from the number of classes N , making it especially useful for the training on datasets with large numbers of persons.

Besides being memory-efficient, Prototype Memory is also more computationally efficient than other softmax-based methods. In Table 5 we compare average computations times of main (and most costly) operations in these methods: cosine similarity calculation between embeddings in the mini-batch and prototypes (weights of classes) in the

TABLE 4. Memory requirements of softmax-based methods, where N is total number of classes in the dataset, M is the size of sampled softmax ($M < N$), and D is the size of embedding.

Method	GPU Memory	Non-GPU Memory
Full Softmax	$D \times N$	-
PPRN [5]	$D \times M$	$D \times N$
D-Softmax-K [28]	$D \times M$	$D \times N$
Prototype Memory	$D \times M$	-

classifier, prototype generation (for Prototype Memory) and RAM-to-GPU transfer of class weights (for other “sampled softmax”-based methods). Computations are performed on 6 NVIDIA GTX 1080 Ti GPUs, with a mini-batch of size 128 and embedding (and also prototype) size of $D = 256$, using parallel acceleration strategy, described in [7]. Total number of classes is set to 1 million ($1M$), numbers of sampled classes are set to $1M$, $500k$, $200k$, $100k$ and $50k$. Number of images per class for Prototype Memory is $k = 4$. Computation times are averaged over multiple iterations (for fair comparisons, for Prototype Memory we start the measurements after the memory is filled to its limits), and given in milliseconds.

TABLE 5. Comparison of average computation times (in milliseconds) of main operations of different softmax-based models. Prototype Memory (PM) is compared to Full Softmax, D-Softmax-K (DS-K) [28] and Positive Plus Randomly Negative (PPRN) [5] models. Numbers are given for cosine similarity scores computation, prototype generation and RAM-to-GPU weights transfer. Values in brackets represent numbers of sampled and total classes. Mini-batch size is 128, prototype size is 256.

Method	Average computation time (ms)		
	Cosine Sim.	Prototypes	RAM-to-GPU
Full Softmax ($1M$)	139	-	-
PPRN ($1M / 1M$) [5]	139	-	107
DS-K ($1M / 1M$) [28]	139	-	107
PM ($1M / 1M$) (our)	139	0.5	-
PPRN ($500k / 1M$)	69	-	54
DS-K ($500k / 1M$)	69	-	54
PM ($500k / 1M$)	69	0.5	-
PPRN ($200k / 1M$)	27	-	25
DS-K ($200k / 1M$)	27	-	25
PM ($200k / 1M$)	27	0.5	-
PPRN ($100k / 1M$)	14	-	12
DS-K ($100k / 1M$)	14	-	12
PM ($100k / 1M$)	14	0.5	-
PPRN ($50k / 1M$)	7	-	7
DS-K ($50k / 1M$)	7	-	7
PM ($50k / 1M$)	7	0.5	-

As we can see from Table 5, “sampled softmax”-based methods are more efficient than full softmax, when the number of sampled classes is small. However, D-Softmax-K and PPRN methods include the costly RAM-to-GPU class weights transfer operation. When the number of sampled classes is larger, it introduces significant additional computational overhead. For some cases this operation could be optimized (for example, when some class weights are already on GPU since previous training iteration), but in general case it still remains as an unavoidable nuisance. On the contrary, Prototype Memory keeps all computations entirely on GPU and is released from this drawback. On the other hand, Prototype Memory introduces the operation of prototype

generation, which also uses some extra computation, but is independent of the number of classes and much faster.

E. EXPERIMENTS WITH MULTI-DOPPELGANGER MINING AND HARDNESS-AWARE EXAMPLE MINING

To evaluate the effectiveness of the proposed Multi-Doppelganger Mining and Hardness-aware example mining algorithms, we have performed experiments with PM-100 model, pre-trained on Glint360k-R for 540,000 iterations using Prototype Memory with $M = 200,000$, $r = 0.2$, $k = 4$ and CosFace with $s = 64$ and $m = 0.4$ as loss function. This model was fine-tuned for 55,000 more iterations, with and without hard example mining methods applied.

For the case of usual training, we used composite mini-batch [56], containing 128 images, sampled with “group-based iterate-and-shuffle” strategy, and 384 images, sampled with “group-based classes-then-images” strategy. We also used $m = 0.5$ as margin value in CosFace.

For the case of hard example mining, we used composite mini-batch, containing 128 images, sampled with “group-based iterate-and-shuffle” strategy, and 384 images, sampled using a combination of Multi-Doppelganger Mining and Hardness-aware example mining, with $h = 0.5$, 12 classes sampled at random and 84 classes sampled using doppelgangers. Margin value $m = 0.5$ was used in CosFace. The results of the experiments are presented in Table 6. Multi-Doppelganger Mining and Hardness-aware example mining methods improve the results on MegaFace benchmark.

TABLE 6. Experiments with Multi-Doppelganger Mining (MDM) and Hardness-aware example mining (HEM).

Method	MegaFace		MegaFace (R)	
	Id	Ver	Id	Ver
Prototype Memory	82.14	96.51	98.34	97.84
Prototype Memory + MDM-HEM	81.88	96.67	98.47	98.37

F. EXPERIMENTS WITH KNOWLEDGE DISTILLATION

We have performed the experiments with Prototype Memory Knowledge Distillation (PMKD). We used embeddings, calculated by EfficientPolyFace [89] network, as teacher embeddings and a ResNet-34 model with memory size of 36,000, $r = 0.2$ and $k = 4$ images per class as a student. Training was performed on Glint360k-M, testing on MegaFace, original and cleaned variants (the latter is more reliable). Results are in Table 7. Proposed knowledge distillation method achieved substantial improvement over the student model.

TABLE 7. Experiments with prototype memory knowledge distillation.

Method	MegaFace		MegaFace (R)	
	Id	Ver	Id	Ver
<i>EPF (Teacher)</i> [89]	83.16	96.94	99.12	98.54
R-34 (Baseline)	81.48	96.02	97.60	98.25
R-34 + PMKD (Student)	81.66	96.94	98.32	98.32

G. COMPARISON WITH STATE-OF-THE-ART

We have compared Prototype Memory to the state-of-the-art models from the literature. To perform fair evaluation [85],

we used strict protocol with removed test / train identity overlaps. For the evaluation we used the PM-100 model, fine-tuned with Multi-Doppelganger Mining and Hardness-aware example mining, as mentioned above.

We also used PM-100 to fine-tune another model, containing Prototype Memory $M = 200,000$, $k = 4$, $r = 0.2$, and using D-Softmax [28] loss function with $D = 0.9$ and $s = 64$. Fine-tuning was performed on Glint-360k-R dataset for 55,000 iterations with mini-batch of 512, containing 128 images, sampled with “group-based iterate-and-shuffle” strategy, and 384 images, sampled with “group-based classes-then-images” strategy. We evaluate both these models on a variety of face recognition benchmarks.

1) CFP-FP, AgeDB-30, CALFW AND CPLFW

We have compared Prototype Memory models to the state-of-the-art on CFP-FP [12] and CPLFW [13] testing datasets, containing large pose variations. We also used AgeDB-30 [14] and CALFW [15] testing datasets, containing large age variations. Results are in Table 8. Models, which explicitly stated, that they have removed train / test identity overlaps (i.e. use strict evaluation protocol), are placed separately from other models, as the training with train / test overlaps could provide unfair and unreliable test results.

For the strict evaluation protocol, Prototype Memory models achieved state-of-the-art results on all four datasets. Prototype Memory model, trained using CosFace with hard example mining, performed better than Prototype Memory model with D-Softmax on all datasets except AgeDB-30, where the latter model achieved slightly superior results.

2) MegaFace

We have compared Prototype Memory models to the state-of-the-art on large-scale MegaFace benchmark, in identification and verification scenarios, original and cleaned versions. Results are presented in Table 9. For the strict evaluation protocol, and among the models, trained using publicly available face datasets (with the exception of WebFace42M, which is a much larger training dataset), Prototype Memory models achieve state-of-the-art results. Prototype Memory model with D-Softmax outperforms Prototype Memory model with CosFace and hard example mining on verification tests and also on identification test for noisy version of MegaFace. On the identification test of cleaned MegaFace, CosFace-based model achieves better results. These results prove the ability of the models, trained with Prototype Memory, to perform accurate face recognition in large-scale scenarios.

Even better results could be achieved with larger training datasets [22], better encoder architectures [115], data augmentation [116] and other methods [117], providing models with the robustness to large age [51] and pose [118] variations and ability to overcome the problems of racial bias [119], domain imbalance [54] and bad image quality [120].

TABLE 8. Comparison with state-of-the-art on CFP-FP, AgeDB-30, CALFW and CPLFW.

Method	CFP-FP	AgeDB-30	CALFW	CPLFW
With possible train / test identity overlap:				
Affinity Loss [90]	96.00	95.90	-	-
Khan et al [91]	97.00	94.40	-	-
D-Softmax [28]	93.07	97.30	-	-
ISL [92]	-	-	92.72	86.13
ShrinkTeaNet-MFNR [93]	95.14	97.63	-	-
RIM [94]	-	-	94.30	90.80
HPDA [95]	-	-	95.90	92.35
SFace [96]	-	-	96.07	93.28
DB [54]	-	97.90	96.08	92.63
MultiFace [97]	-	-	96.08	93.32
QAMFace [98]	-	-	96.10	92.85
OL [99]	-	-	96.12	92.65
LATSE [9]	-	-	96.20	93.48
AFRN [100]	95.56	95.35	96.30	93.48
SRANet [101]	95.60	98.47	-	-
ACNN [102]	95.85	98.57	-	-
MBFN [103]	96.75	98.96	-	-
CurricularFace [32]	98.37	98.32	96.20	93.13
ProbFace [104]	98.41	97.90	96.02	93.53
MagFace [36]	98.46	98.17	96.15	92.87
DiscFace [105]	98.54	98.35	96.15	92.87
GroupFace [35]	98.63	98.28	96.20	93.17
BroadFace [34]	98.63	98.38	96.20	93.17
DAIL [106]	98.70	98.20	-	-
DDL [107]	99.06	-	-	94.20
Sub-center ArcFace [8]	99.11	98.35	-	-
Partial FC (r=0.1) [5]	99.33	98.57	96.20	94.83
Partial FC (r=1.0) [5]	99.33	98.55	96.21	94.78
BioMetricNet [108]	99.35	96.12	97.07	95.60
RetinaFace + ArcFace [109]	99.49	98.60	-	-
Without train / test identity overlap:				
Co-Mining (MsCeleb-R) [110]	93.32	95.80	93.28	85.70
MTLFace [111]	-	96.23	95.62	-
NPCFace [33]	95.09	97.77	95.60	89.42
SST (AM-Softmax) [112]	95.10	97.20	94.62	88.35
MV-AM-Softmax-a (0.2) [47]	95.30	98.00	95.63	89.19
Search-Softmax [113]	95.64	97.75	95.40	89.50
MV-AM-Softmax-f (0.25) [47]	95.70	98.05	95.45	89.69
Co-Mining (VGGFace2-R) [110]	95.87	94.05	91.06	87.31
LMFA + TDN [114]	96.86	97.92	-	-
MTCNN + ArcFace [7]	98.37	98.15	95.45	92.08
Prototype Memory + CosFace + MDM-HEM	99.23	98.28	95.97	94.98
Prototype Memory + D-Softmax	99.19	98.37	95.95	94.72

V. CONCLUSION

In this paper we proposed Prototype Memory - a novel face representation learning model, opening the possibility to train state-of-the-art face recognition architectures on the datasets of any size. Prototype Memory consists of the memory module for storing class prototypes, and algorithms to perform operations with it. Prototypes are generated online, using exemplar embeddings, presented in the current mini-batch. New prototypes are enqueued to the memory. Prototypes in memory are refreshed and kept up-to-date with the state of the encoder. When the memory is filled, the oldest prototypes are dequeued and disposed of. Prototypes in memory are

used in the role of classifier weights for softmax-based face representation learning. Prototype memory is useful for preventing the problem of “prototype obsolescence”. Like other “sampled softmax”-based models, it is computationally and memory-efficient. We have performed extensive experimental evaluations on popular face recognition benchmarks and proved the effectiveness of the proposed model. We have compared the performance of Prototype Memory to other “sampled softmax”-based models and demonstrated its superiority both in terms of face recognition accuracy and memory efficiency. We proposed Multi-Doppelganger Mining and Hardness-aware example mining methods to improve the

TABLE 9. Evaluation on MegaFace benchmark, where † denotes the usage of large-scale private training datasets.

Method	MegaFace		MegaFace (R)	
	Id	Ver	Id	Ver
With possible train / test identity overlap:				
SphereFace [23]	72.73	85.56	-	-
SphereFace+ [121]	73.03	-	-	-
CCL with AAM [122]	73.74	-	-	-
AM-Softmax [46]	75.23	87.06	-	-
Ring Loss [123]	75.22	-	-	-
EEL [83]	76.40	90.70	-	-
ShuffleFaceNet [124]	77.40	93.00	-	-
Fair Loss [55]	77.45	92.87	-	-
Wang et al [125]	77.74	79.24	-	-
BUPT-CBFace [126]	79.57	95.20	-	-
CLMLE [127]	79.68	91.85	-	-
UniformFace [48]	79.98	95.36	-	-
Marginal Loss [128]	80.28	92.64	-	-
m-LinCos-Softmax [129]	-	-	88.69	91.22
AdaptiveFace [29]	-	-	95.02	95.61
DB [54]	-	-	96.35	96.56
ATAM [130]	-	-	96.51	97.14
D-Softmax [28]	-	-	97.02	-
NPT-Loss [131]	-	-	97.67	98.54
MBCConv [132]	80.29	96.47	97.80	97.94
AirFace [133]	80.80	96.52	98.04	97.93
DUL _{r_{rgs}} [134]	-	-	98.12	98.78
SRANet [101]	-	-	98.34	98.38
ACNN [102]	-	-	98.35	98.42
OL [99]	81.05	97.16	98.37	98.80
IBN-block [135]	-	-	98.43	-
Circle Loss [136]	-	-	98.50	-
SFace [96]	81.15	97.11	98.50	98.61
IDA-R [37]	-	-	98.56	-
DUL _{cls} [134]	-	-	98.60	98.67
BA [137]	-	-	98.62	-
MultiFace [97]	80.91	97.37	98.62	98.85
DiscFace [105]	81.23	97.44	-	-
BroadFace [34]	81.33	97.56	98.70	98.95
CurricularFace [32]	81.26	97.26	98.71	98.64
GroupFace [35]	81.31	97.35	98.74	98.79
DFN-152 [52]	82.11	-	-	-
IMDb-Face [26]	84.06	-	-	-
CDP† [138]	84.74	-	-	-
CosFace + DAG† [139]	85.62	97.26	-	-
Sub-center ArcFace [8]	-	-	98.78	98.69
Partial FC (r=0.1) [5]	-	-	98.94	99.10
Partial FC (r=1.0) [5]	-	-	99.13	98.98
LATSE [9]	-	-	99.14	99.19
Without train / test identity overlap:				
Center Loss [43]	65.23	76.52	-	-
AM-Softmax [46]	72.47	84.44	-	-
LightCNN-29 [140]	73.75	85.13	-	-
RegularFace [49]	75.61	91.13	-	-
COCO [141]	76.57	-	-	-
MTLFace [111]	77.06	-	-	-
FAN-Face [142]	78.32	92.83	-	-
DS-Res64 [143]	80.81	-	-	-
SST (AM-Softmax) [112]	-	-	96.27	96.96
Search-Softmax [113]	-	-	96.97	97.84
NPCFace [33]	-	-	97.75	98.07
CosFace [7]	80.56	96.56	97.91	97.91
MV-AM-Softmax-a [47]	-	-	98.00	98.31
ArcFace [7]	81.03	96.98	98.35	98.48
LMFA + TDN [114]	81.42	95.22	-	-
CosFace-Private† [6]	82.72	96.65	-	-
MC - MegaFace2† [20]	-	-	98.96	98.97
MC - FaceGraph† [20]	-	-	99.02	98.94
WebFace42M [22]	-	-	99.11	-
PM + CosFace + MDM-HEM	81.88	96.67	98.47	98.37
PM + D-Softmax	82.26	97.39	98.38	98.72

training of models, based on Prototype Memory, with the help of the generation of more informative mini-batches. We also described a knowledge distillation method, suitable for using with Prototype Memory, and proved its effectiveness.

REFERENCES

- [1] M. Wang and W. Deng, "Deep face recognition: A survey," *Neurocomputing*, vol. 429, pp. 215–244, Mar. 2021.
- [2] H. Du, H. Shi, D. Zeng, X.-P. Zhang, and T. Mei, "The elements of end-to-end deep face recognition: A survey of recent advances," 2020, *arXiv:2009.13290*.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, Feb. 2015.
- [4] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "MS-Celeb-1M: A dataset and benchmark for large-scale face recognition," in *Proc. ECCV*, 2016, pp. 87–102.
- [5] X. An, X. Zhu, Y. Xiao, L. Wu, M. Zhang, Y. Gao, B. Qin, D. Zhang, and Y. Fu, "Partial FC: Training 10 million identities on a single machine," 2020, *arXiv:2010.05222*.
- [6] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "CosFace: Large margin cosine loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5265–5274.
- [7] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4690–4699.
- [8] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou, "Sub-center ArcFace: Boosting face recognition by large-scale noisy web faces," in *Proc. ECCV*. Glasgow, U.K.: Springer, 2020, pp. 741–757.
- [9] Y. Huang, S. Qiu, W. Zhang, X. Luo, and J. Wang, "More information supervised probabilistic deep face embedding learning," in *Proc. ICML*, 2020, pp. 4484–4494.
- [10] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1701–1708.
- [11] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard, "The MegaFace benchmark: 1 million faces for recognition at scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 4873–4882.
- [12] S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs, "Frontal to profile face verification in the wild," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2016, pp. 1–9.
- [13] T. Zheng and W. Deng, "Cross-pose LFW: A database for studying cross-pose face recognition in unconstrained environments," Beijing Univ. Posts Telecommun., Beijing, China, Tech. Rep. 18-01, 2018.
- [14] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou, "AgeDB: The first manually collected, in-the-wild age database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 51–59.
- [15] T. Zheng, W. Deng, and J. Hu, "Cross-age LFW: A database for studying Cross-age face recognition in unconstrained environments," 2017, *arXiv:1708.08197*.
- [16] M. Wang, W. Deng, J. Hu, X. Tao, and Y. Huang, "Racial faces in the wild: Reducing racial bias by information maximization adaptation network," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 692–702.
- [17] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, J. Cheney, and P. Grother, "IARPA Janus benchmark-C: Face dataset and protocol," in *Proc. Int. Conf. Biometrics (ICB)*, Feb. 2018, pp. 158–165.
- [18] M. Singh, R. Singh, M. Vatsa, N. K. Ratha, and R. Chellappa, "Recognizing disguised faces in the wild," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 1, no. 2, pp. 97–108, Apr. 2019.
- [19] N. Damer, F. Boutros, M. Süßmilch, M. Fang, F. Kirchbuchner, and A. Kuijper, "Masked face recognition: Human vs. machine," 2021, *arXiv:2103.01924*.
- [20] Y. Zhang, W. Deng, M. Wang, J. Hu, X. Li, D. Zhao, and D. Wen, "Global-local GCN: Large-scale label noise cleansing for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 7731–7740.
- [21] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 815–823.
- [22] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du, and J. Zhou, "WebFace260M: A benchmark unveiling the power of million-scale deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 10492–10502.

- [23] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "SphereFace: Deep hypersphere embedding for face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 212–220.
- [24] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *Proc. 13th IEEE Int. Conf. Autom. Face Gesture Recognit. (FG)*, May 2018, pp. 67–94.
- [25] A. Bansal, A. Nanduri, C. D. Castillo, R. Ranjan, and R. Chellappa, "UMDFaces: An annotated face dataset for training deep networks," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2017, pp. 464–473.
- [26] F. Wang, L. Chen, C. Li, S. Huang, Y. Chen, C. Qian, and C. C. Loy, "The devil of face recognition is in the noise," in *Proc. ECCV*, Sep. 2018, pp. 765–780.
- [27] X. Zhang, L. Yang, J. Yan, and D. Lin, "Accelerated training for massive classification via dynamic class selection," in *Proc. AAAI*, 2018, pp. 7566–7573.
- [28] L. He, Z. Wang, Y. Li, and S. Wang, "Softmax dissection: Towards understanding intra- and inter-class objective for embedding learning," in *Proc. AAAI*, 2020, pp. 10957–10964.
- [29] H. Liu, X. Zhu, Z. Lei, and S. Z. Li, "AdaptiveFace: Adaptive margin and sampling for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 11947–11956.
- [30] E. Smirnov, A. Melnikov, S. Novoselov, E. Lukanets, and G. Lavrentyeva, "Doppelganger mining for face representation learning," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops (ICCVW)*, Oct. 2017, pp. 1916–1923.
- [31] E. Smirnov, E. Ivanova, A. Melnikov, I. Kalinovskiy, A. Oleinik, and E. Lukanets, "Hard example mining with auxiliary embeddings," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 37–46.
- [32] Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang, "CurricularFace: Adaptive curriculum learning loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5901–5910.
- [33] D. Zeng, H. Shi, H. Du, J. Wang, Z. Lei, and T. Mei, "NPCFace: Negative-positive collaborative training for large-scale face recognition," 2020, *arXiv:2007.10172*.
- [34] Y. Kim, W. Park, and J. Shin, "BroadFace: Looking at tens of thousands of people at once for face recognition," in *Proc. ECCV. Glasgow, U.K.: Springer*, 2020, pp. 536–552.
- [35] Y. Kim, W. Park, M.-C. Roh, and J. Shin, "GroupFace: Learning latent groups and constructing group-based representations for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5621–5630.
- [36] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A universal representation for face recognition and quality assessment," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 14225–14234.
- [37] J. Liu, Y. Wu, Y. Wu, Z. Li, C. Ken, D. Liang, and J. Yan, "Inter-class discrepancy alignment for face recognition," 2021, *arXiv:2103.01559*.
- [38] X. Xu, Q. Meng, Y. Qin, J. Guo, C. Zhao, F. Zhou, and Z. Lei, "Searching for alignment in face recognition," in *Proc. AAAI*, 2021, pp. 3065–3073.
- [39] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *Proc. BMVC*, 2015, pp. 1–12.
- [40] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, Oct. 2016.
- [41] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-shot multi-level face localisation in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5203–5212.
- [42] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2005, pp. 539–546.
- [43] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *Proc. ECCV*, 2016, pp. 499–515.
- [44] R. Ranjan, C. D. Castillo, and R. Chellappa, "L₂-constrained softmax loss for discriminative face verification," 2017, *arXiv:1703.09507*.
- [45] F. Wang, X. Xiang, J. Cheng, and A. L. Yuille, "NormFace: L₂ hypersphere embedding for face verification," in *Proc. 25th ACM Int. Conf. Multimedia*, Oct. 2017, pp. 1041–1049.
- [46] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Process. Lett.*, vol. 25, no. 7, pp. 926–930, Jul. 2018.
- [47] X. Wang, S. Zhang, S. Wang, T. Fu, H. Shi, and T. Mei, "Mis-classified vector guided softmax loss for face recognition," in *Proc. AAAI*, 2020, pp. 12241–12248.
- [48] Y. Duan, J. Lu, and J. Zhou, "UniformFace: Learning deep equidistributed representation for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3415–3424.
- [49] K. Zhao, J. Xu, and M.-M. Cheng, "RegularFace: Deep face recognition via exclusive regularization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 1136–1144.
- [50] S. S. Khalid, M. Awais, Z.-H. Feng, C.-H. Chan, A. Farooq, A. Akbari, and J. Kittler, "Resolution invariant face recognition using a distillation approach," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 2, no. 4, pp. 410–420, Oct. 2020.
- [51] J. Zhao, S. Yan, and J. Feng, "Towards age-invariant face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 1, pp. 474–487, Jan. 2022.
- [52] M. He, J. Zhang, S. Shan, M. Kan, and X. Chen, "Deformable face net for pose invariant face recognition," *Pattern Recognit.*, vol. 100, Apr. 2020, Art. no. 107113.
- [53] M. Wang and W. Deng, "Mitigating bias in face recognition using skewness-aware reinforcement learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9322–9331.
- [54] D. Cao, X. Zhu, X. Huang, J. Guo, and Z. Lei, "Domain balancing: Face recognition on long-tailed domains," in *Proc. CVPR*, Jun. 2020, pp. 5671–5679.
- [55] B. Liu, W. Deng, Y. Zhong, M. Wang, J. Hu, X. Tao, and Y. Huang, "Fair loss: Margin-aware reinforcement learning for deep face recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 10052–10061.
- [56] E. Smirnov, A. Oleinik, A. Lavrentev, E. Shulga, V. Galyuk, N. Garaev, M. Zakuanova, and A. Melnikov, "Face representation learning using composite mini-batches," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 551–559.
- [57] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "Self-restrained triplet loss for accurate masked face recognition," 2021, *arXiv:2103.01716*.
- [58] M. Milakov and N. Gimelshein, "Online normalizer calculation for softmax," 2018, *arXiv:1805.02867*.
- [59] M. Dukhan and A. Ablavatski, "Two-pass softmax algorithm," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. Workshops (IPDPSW)*, May 2020, pp. 386–395.
- [60] R. Bamler and S. Mandt, "Extreme classification via adversarial softmax approximation," in *Proc. ICLR*, 2020. [Online]. Available: <https://dblp.uni-trier.de/rec/conf/iclr/BamlerM20.html?view=bibtex>
- [61] K. Shim, M. Lee, I. Choi, Y. Boo, and W. Sung, "SVD-Softmax: Fast softmax approximation on large vocabulary neural networks," in *Proc. NeurIPS*, 2017, pp. 5469–5479.
- [62] E. Grave, A. Joulin, M. Cissé, D. Grangier, and H. Jégou, "Efficient softmax approximation for GPUs," in *Proc. ICML*, 2017, pp. 1302–1310.
- [63] Y. Wu, J. Li, Y. Kong, and Y. Fu, "Deep convolutional neural network with independent softmax for large scale face recognition," in *Proc. 24th ACM Int. Conf. Multimedia*, Oct. 2016, pp. 1063–1067.
- [64] M. Wydmuch, K. Jasinska, M. Kuznetsov, R. Busa-Fekete, and K. Dembczynski, "A no-regret generalization of hierarchical softmax to extreme multi-label classification," in *Proc. NeurIPS*, 2018, pp. 6358–6368.
- [65] Y. Shen, S. Tan, C. Pal, and A. Courville, "Self-organized hierarchical softmax," 2017, *arXiv:1707.08588*.
- [66] S. J. Reddi, S. Kale, F. Yu, D. Holtmann-Rice, J. Chen, and S. Kumar, "Stochastic negative mining for learning with large output spaces," in *Proc. AISTATS*, 2019, pp. 1940–1949.
- [67] A. S. Rawat, J. Chen, F. X. X. Yu, A. T. Suresh, and S. Kumar, "Sampled softmax with random Fourier features," in *Proc. NeurIPS*, 2019, pp. 13857–13867.
- [68] G. Blanc and S. Rendle, "Adaptive sampled softmax with kernel based sampling," in *Proc. ICML*, 2018, pp. 590–599.
- [69] Y. Bai, S. Goldman, and L. Zhang, "TAPAS: Two-pass approximate adaptive sampling for softmax," 2017, *arXiv:1707.03073*.
- [70] X. Zhu, H. Liu, Z. Lei, H. Shi, F. Yang, D. Yi, G. Qi, and S. Z. Li, "Large-scale bisample learning on ID versus spot face recognition," 2018, *arXiv:1806.03018*.
- [71] H. Qi, M. Brown, and D. G. Lowe, "Low-shot learning with imprinted weights," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5822–5830.

- [72] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. NeurIPS*, 2017, pp. 4080–4090.
- [73] X. Zhang, R. Zhao, Y. Qiao, and H. Li, "RBF-Softmax: Learning deep representative prototypes with radial basis function softmax," in *Proc. ECCV*, 2020, pp. 296–311.
- [74] J. Li, P. Zhou, C. Xiong, and S. C. Hoi, "Prototypical contrastive learning of unsupervised representations," in *Proc. ICLR*, 2021. [Online]. Available: <https://dblp.uni-trier.de/rec/conf/iclr/0001ZXH21.html?view=bibtex>
- [75] M. De Lange and T. Tuytelaars, "Continual prototype evolution: Learning online from non-stationary data streams," in *Proc. ICCV*, 2021, pp. 8250–8259.
- [76] Z. Zhong, L. Zheng, Z. Luo, S. Li, and Y. Yang, "Invariance matters: Exemplar memory for domain adaptive person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 598–607.
- [77] X. Wang, H. Zhang, W. Huang, and M. R. Scott, "Cross-batch memory for embedding learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 6388–6397.
- [78] J. Liu, J. Zhang, Y. Yang, W. Li, C. Zhang, and Y. Sun, "Memory-based jitter: Improving visual recognition on long-tailed data with diversity in memory," 2020, *arXiv:2008.09809*.
- [79] Q. Wang, Y. Ma, K. Zhao, and Y. Tian, "A comprehensive survey of loss functions in machine learning," *Ann. Data Sci.*, pp. 1–26, Apr. 2020. [Online]. Available: <https://link.springer.com/article/10.1007/s40745-020-00253-5>
- [80] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*.
- [81] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *Int. J. Comput. Vis.*, vol. 129, no. 6, pp. 1789–1819, Jun. 2021.
- [82] D. Svitov and S. Alyamkin, "MarginDistillation: Distillation for margin-based softmax," 2020, *arXiv:2003.02586*.
- [83] X. Wu, R. He, Y. Hu, and Z. Sun, "Learning an evolutionary embedding via massive knowledge distillation," *Int. J. Comput. Vis.*, vol. 128, nos. 8–9, pp. 2089–2106, Sep. 2020.
- [84] D. Nekhaev, S. Milyaev, and I. Laptev, "Margin based knowledge distillation for mobile face recognition," in *Proc. 12th Int. Conf. Mach. Vis. (ICMV)*, Jan. 2020, Art. no. 1143300.
- [85] C. Dulhanty and A. Wong, "Investigating the impact of inclusion in face recognition training data on individual face identification," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, Feb. 2020, pp. 244–250.
- [86] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Proc. Workshop Faces Real-Life Images, Detection, Alignment, Recognit.*, 2008, pp. 1–11.
- [87] *DeepGlint*. Accessed: Dec. 22, 2020. [Online]. Available: <http://trillionpairs.deepglint.com>
- [88] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [89] Y. Liu, G. Song, M. Zhang, J. Liu, Y. Zhou, and J. Yan, "Towards flops-constrained face recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 2698–2702.
- [90] M. Hayat, S. Khan, S. W. Zamir, J. Shen, and L. Shao, "Gaussian affinity for max-margin class imbalanced learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6469–6479.
- [91] S. Khan, M. Hayat, S. W. Zamir, J. Shen, and L. Shao, "Striking the right balance with uncertainty," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 103–112.
- [92] J. Zhou, X. Jia, L. Shen, Z. Wen, and Z. Ming, "Improved softmax loss for deep learning-based face and expression recognition," *Cognit. Comput. Syst.*, vol. 1, no. 4, pp. 97–102, Dec. 2019.
- [93] C. N. Duong, K. Luu, K. G. Quach, and N. Le, "ShrinkTeaNet: Million-scale lightweight face recognition via shrinking teacher-student networks," 2019, *arXiv:1905.10620*.
- [94] H. Kong, J. Zhao, X. Tu, J. Xing, S. Shen, and J. Feng, "Cross-resolution face recognition via prior-aided face hallucination and residual knowledge distillation," 2019, *arXiv:1905.10777*.
- [95] Q. Wang, T. Wu, H. Zheng, and G. Guo, "Hierarchical pyramid diverse attention networks for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8326–8335.
- [96] Y. Zhong, W. Deng, J. Hu, D. Zhao, X. Li, and D. Wen, "SFace: Sigmoid-constrained hypersphere loss for robust face recognition," *IEEE Trans. Image Process.*, vol. 30, pp. 2587–2598, 2021.
- [97] J. Xu, T. Guo, Y. Xu, Z. Xu, and K. Bai, "MultiFace: A generic training mechanism for boosting face recognition performance," *Neurocomputing*, vol. 448, pp. 40–47, Aug. 2021.
- [98] H. Zhao, Y. Shi, X. Tong, X. Ying, and H. Zha, "Qamface: Quadratic additive angular margin loss for face recognition," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 1901–1905.
- [99] S. Yang, W. Deng, M. Wang, J. Du, and J. Hu, "Orthogonality loss: Learning discriminative representations for face recognition," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 6, pp. 2301–2314, Jun. 2021.
- [100] B.-N. Kang, Y. Kim, B. Jun, and D. Kim, "Attentional feature-pair relation networks for accurate face recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 5472–5481.
- [101] H. Ling, J. Wu, L. Wu, J. Huang, J. Chen, and P. Li, "Self residual attention network for deep face recognition," *IEEE Access*, vol. 7, pp. 55159–55168, 2019.
- [102] H. Ling, J. Wu, J. Huang, J. Chen, and P. Li, "Attention-based convolutional neural network for deep face recognition," *Multimedia Tools Appl.*, vol. 79, nos. 9–10, pp. 5595–5616, Mar. 2020.
- [103] J. Zhang, B. Ye, and X. Luo, "MBFN: A multi-branch face network for facial analysis," in *Proc. 2nd Int. Conf. Algorithms, Comput. Artif. Intell.*, Dec. 2019, pp. 542–549.
- [104] K. Chen, Q. Lv, and T. Yi, "Fast and reliable probabilistic face embeddings in the wild," 2021, *arXiv:2102.04075*.
- [105] I. Kim, S. Han, S.-J. Park, J.-W. Baek, J. Shin, J.-J. Han, and C. Choi, "DiscFace: Minimum discrepancy learning for deep face recognition," in *Proc. ACCV*, 2020, pp. 358–374.
- [106] G. Wang, L. Chen, T. Liu, M. He, and J. Luo, "DAIL: Dataset-aware and invariant learning for face recognition," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 8172–8179.
- [107] Y. Huang, P. Shen, Y. Tai, S. Li, X. Liu, J. Li, F. Huang, and R. Ji, "Improving face recognition from hard samples via distribution distillation loss," in *Proc. ECCV*, Aug. 2020, pp. 138–154.
- [108] A. Ali, M. Testa, T. Bianchi, and E. Magli, "BioMetricNet: Deep unconstrained face verification through learning of metrics regularized onto Gaussian distributions," in *Proc. ECCV*. Glasgow, U.K.: Springer, 2020, pp. 133–149.
- [109] J. Deng, J. Guo, Y. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-stage dense face localisation in the wild," 2019, *arXiv:1905.00641*.
- [110] X. Wang, S. Wang, H. Shi, J. Wang, and T. Mei, "Co-mining: Deep face recognition with noisy labels," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9358–9367.
- [111] Z. Huang, J. Zhang, and H. Shan, "When age-invariant face recognition meets face age synthesis: A multi-task learning framework," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 7282–7291.
- [112] H. Du, H. Shi, Y. Liu, J. Wang, Z. Lei, D. Zeng, and T. Mei, "Semi-Siamese training for shallow face learning," in *Proc. ECCV*. Glasgow, U.K.: Springer, 2020, pp. 36–53.
- [113] X. Wang, S. Wang, C. Chi, S. Zhang, and T. Mei, "Loss function search for face recognition," in *Proc. ICML*, 2020, pp. 10029–10038.
- [114] P. Wang, F. Su, Z. Zhao, Y. Guo, Y. Zhao, and B. Zhuang, "Deep class-skewed learning for face recognition," *Neurocomputing*, vol. 363, pp. 35–45, Oct. 2019.
- [115] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, "Transformers in vision: A survey," 2021, *arXiv:2101.01169*.
- [116] A. Buslaev, V. I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, and A. A. Kalinin, "Albumentations: Fast and flexible image augmentations," *Information*, vol. 11, no. 2, p. 125, Feb. 2020.
- [117] A. Sanakoyev, P. Ma, V. Tschernezki, and B. Ommer, "Improving deep metric learning by divide and conquer," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Sep. 16, 2021, doi: [10.1109/TPAMI.2021.3113270](https://doi.org/10.1109/TPAMI.2021.3113270).
- [118] S. B. Ahmed, S. F. Ali, J. Ahmad, M. Adnan, and M. M. Fraz, "On the frontiers of pose invariant face recognition: A review," *Artif. Intell. Rev.*, vol. 53, no. 4, pp. 2571–2634, Apr. 2020.
- [119] T. Sixta, J. C. J. Junior, P. Buch-Cardona, E. Vazquez, and S. Escalera, "FairFace challenge at ECCV 2020: Analyzing bias in face recognition," in *Proc. ECCV*, Aug. 2020, pp. 463–481.
- [120] L. S. Luevano, L. Chang, H. Heydi Mendez-Vazquez, Y. Martinez-Diaz, and M. Gonzalez-Mendoza, "A study on the performance of unconstrained very low resolution face recognition: Analyzing current trends and new research directions," *IEEE Access*, vol. 9, pp. 75470–75493, 2021.
- [121] J. Liu, S. J. Gibson, and M. Osadchy, "Learning to support: Exploiting structure information in support sets for one-shot learning," 2018, *arXiv:1808.07270*.
- [122] X. Qi and L. Zhang, "Face recognition via centralized coordinate learning," 2018, *arXiv:1801.05678*.
- [123] Y. Zheng, D. K. Pal, and M. Savvides, "Ring loss: Convex feature normalization for face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5089–5097.

- [124] Y. Martinez-Diaz, L. S. Luevano, H. Mendez-Vazquez, M. Nicolas-Diaz, L. Chang, and M. Gonzalez-Mendoza, "ShuffleFaceNet: A lightweight face architecture for efficient and highly-accurate face recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 2712–2728.
- [125] Z. Wang, K. He, Y. Fu, R. Feng, Y.-G. Jiang, and X. Xue, "Multi-task deep neural network for joint face recognition and facial attribute prediction," in *Proc. ACM Int. Conf. Multimedia Retr.*, Jun. 2017, pp. 365–374.
- [126] Y. Zhang and W. Deng, "Class-balanced training for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 824–825.
- [127] C. Huang, Y. Li, C. C. Loy, and X. Tang, "Deep imbalanced learning for face recognition and attribute prediction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 11, pp. 2781–2794, Nov. 2020.
- [128] J. Deng, Y. Zhou, and S. Zafeiriou, "Marginal loss for deep face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 60–68.
- [129] W.-F. Ou, L.-M. Po, C. Zhou, Y.-J. Zhang, L.-T. Feng, Y. A. U. Rehman, and Y.-Z. Zhao, "LinCos-softmax: Learning angle-discriminative face representations with linearity-enhanced cosine logits," *IEEE Access*, vol. 8, pp. 109758–109769, 2020.
- [130] S. M. Iranmanesh, A. Dabouei, and N. M. Nasrabadi, "Attribute adaptive margin softmax loss using privileged information," in *Proc. BMVC*, 2020, pp. 1–13.
- [131] S. S. Khalid, M. Awais, C.-H. Chan, Z. Feng, A. Farooq, A. Akbari, and J. Kittler, "NPT-loss: A metric loss with implicit mining for face recognition," 2021, *arXiv:2103.03503*.
- [132] Y. Lyu, J. Jiang, K. Zhang, Y. Hua, and M. Cheng, "Factorizing and reconstituting large-kernel MBConv for lightweight face recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 2689–2697.
- [133] X. Li, F. Wang, Q. Hu, and C. Leng, "AirFace: Lightweight and efficient model for face recognition," in *Proc. ICCV Workshops*, Oct. 2019, pp. 2678–2682.
- [134] J. Chang, Z. Lan, C. Cheng, and Y. Wei, "Data uncertainty learning in face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 5710–5719.
- [135] Y. Qing, Y. Zhao, Y. Shi, D. Chen, Y. Lin, and Y. Peng, "Improve cross-domain face recognition with IBN-block," in *Proc. Big Data*, Dec. 2018, pp. 4613–4618.
- [136] Y. Sun, C. Cheng, Y. Zhang, C. Zhang, L. Zheng, Z. Wang, and Y. Wei, "Circle loss: A unified perspective of pair similarity optimization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 6398–6407.
- [137] H. Wei, P. Lu, and Y. Wei, "Balanced alignment for face recognition: A joint learning approach," 2020, *arXiv:2003.10168*.
- [138] X. Zhan, Z. Liu, J. Yan, D. Lin, and C. C. Loy, "Consensus-driven propagation in massive unlabeled data for face recognition," in *Proc. ECCV*, Sep. 2018, pp. 568–583.
- [139] A. Dabouei, F. Taherkhani, S. Soleymani, J. Dawson, and N. M. Nasrabadi, "Boosting deep face recognition via disentangling appearance and geometry," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2020, pp. 320–329.
- [140] X. Wu, R. He, Z. Sun, and T. Tan, "A light CNN for deep face representation with noisy labels," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 11, pp. 2884–2896, Nov. 2018.
- [141] Y. Liu, H. Li, and X. Wang, "Rethinking feature discrimination and polymerization for large-scale recognition," 2017, *arXiv:1710.00870*.
- [142] J. Yang, A. Bulat, and G. Tzimiropoulos, "FAN-Face: A simple orthogonal improvement to deep face recognition," in *Proc. AAAI*, 2020, pp. 12621–12628.
- [143] Z. Liu, W. Hong, H. Zhang, and J. Ma, "Face recognition with dense supervision," *Neurocomputing*, vol. 387, pp. 100–109, Apr. 2020.



EVGENY SMIRNOV received the Specialist degree in information systems and technology from ITMO University, Saint Petersburg, Russia, in 2011. He is currently a Senior Researcher with the Speech Technology Center (STC). His research interests include computer vision, face recognition, and representation learning.



NIKITA GARAEV received the B.E. degree from Saint Petersburg Electrotechnical University (ETU "LETI"), Russia, in 2019, and the M.S. degree from ITMO University, Russia, in 2021. He is currently working with the Speech Technology Center. His research interests include computer vision and face recognition.



VASILII GALYUK received the B.Sc. and master's degrees from ITMO University, Saint Petersburg, Russia, in 2015 and 2017, respectively. He is currently a Researcher with the Speech Technology Center (STC). His research interests include computer vision, image classification, and face detection.



EVGENY LUKYANETS received the master's degree in information systems and technology from ITMO University, Saint Petersburg, Russia, in 2015. He is currently a Developer with ITMO University and the Technical Product Lead with the Speech Technology Center (STC). His research interests include computer vision and face recognition.

...