

FaceCoresetNet: Differentiable Coresets for Face Set Recognition

Gil Shapira^{1,2} and Yosi Keller¹

¹ Bar-Ilan University, Ramat Gan, Israel

² Samsung Semiconductor Israel R&D Center (SIRC)
{ggiillsshhaappiirraa, yosi.keller}@gmail.com

Abstract

In set-based face recognition, we aim to compute the most discriminative descriptor from an unbounded set of images and videos showing a single person. A discriminative descriptor balances two policies when aggregating information from a given set. The first is a quality-based policy: emphasizing high-quality and down-weighting low-quality images. The second is a diversity-based policy: emphasizing unique images in the set and down-weighting multiple occurrences of similar images as found in video clips which can overwhelm the set representation. This work frames face-set representation as a differentiable coreset selection problem. Our model learns how to select a small coreset of the input set that balances quality and diversity policies using a learned metric parameterized by the face quality, optimized end-to-end. The selection process is a differentiable farthest-point sampling (FPS) realized by approximating the non-differentiable Argmax operation with differentiable sampling from the Gumbel-Softmax distribution of distances. The small coreset is later used as queries in a self and cross-attention architecture to enrich the descriptor with information from the whole set. Our model is order-invariant and linear in the input set size. We set a new SOTA to set face verification on the IJB-B and IJB-C datasets. Our code is publicly available at <https://github.com/ligaripash/FaceCoresetNet>.

Introduction

Face Set Recognition (FSR) is a variant of Face Recognition (FR) that involves recognizing an individual's identity based on a *template*, that is an unbounded and unordered set of images and video clips depicting a particular face instead of a single photograph. A template is shown in the Fig. 1. Given the added data, FSR is more accurate compared to FR. The increased accuracy and proliferation of face media available on the Internet and public cameras make FSR an attractive solution for 'in the wild' situations, such as in unconstrained surveillance and Homeland Security scenarios (Kalka et al. 2018) where image quality is lacking due to low resolution, extreme pose, illumination, occlusions, and recognition of a single image is difficult.

FSR consists of verification, which means confirming the identity of an input query template (also referred to as a

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

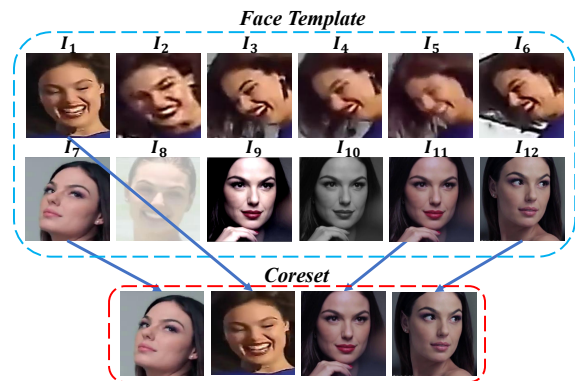


Figure 1: Template coreset selection. in the round blue rectangle, an input face set with 12 images containing a video clip sample (I_1 to I_6) and six photographs of the same person (I_7 to I_{12}). The coreset selection process balances quality and diversity. The final coreset in the red round rectangle contains four high-quality images with diverse poses and expressions.

probe) and identification, which entails searching a database of registered subjects (known as a gallery) to identify an unknown probe person. FSR schemes encode the images in a template into a fixed-size vector descriptor to allow computationally efficient similarity search. As the templates consist of an arbitrary number of images and video frames, the image aggregation process has to be efficient with respect to the template's size. Furthermore, since the template is an unordered set of images, the aggregation should be permutation invariant. Additionally, the aggregation should prioritize informative images while downplaying uninformative, low-quality (blurry, poorly lighted) images. Finally, the process of feature fusion should factor in "face burstiness". If a template includes a burst of similar frames as in video clips, the burst should not overpower the template descriptor.

Some aggregation schemes apply average pooling (Chen et al. 2015; Parkhi, Vedaldi, and Zisserman 2015) to the set of features. However, this method presents a significant drawback. All images within the set are weighted equally, regardless of their quality. To address this issue, (Kim, Jain, and Liu 2022; Meng et al. 2021) suggest facial features

whose magnitude is proportional to image quality. To improve the aggregation, other works model intra-set relationships using reinforcement learning (Liu et al. 2018), recurrent models (Gong, Shi, and Jain 2019), or self-attention (Gong et al. 2019; Liu et al. 2019; Liu, Yan, and Ouyang 2017; Xie and Zisserman 2018). In contrast to intra-set modeling, face burstiness received only limited attention in the literature (Wang, Zhao, and Wu 2022).

While Recurrent methods (Graves and Graves 2012; Gong, Shi, and Jain 2019) are effective in sequence modeling, they are unsuitable for modeling orderless sets as they are not permutation invariant. On the other hand, Multihead Self Attention (MSA) (Vaswani et al. 2017; Liu et al. 2018) is an effective permutation-invariant approach to model self-attention among template features, but it becomes impractical for large templates since it is quadratic in the template size.

In this work, we propose a feature aggregation approach, denoted *FaceCoresetNet*, that selects a small coreset we dub *Core-Template* of the face template with balanced quality and diversity. Given N data points, a coreset is a smaller subset of $K \ll N$ sampled points that approximate the original dataset with respect to a given metric. It is often used to speed up training or reduce the memory and storage requirements of large datasets (Mirzasoaleiman, Bilmes, and Leskovec 2020; Campbell and Broderick 2018; Pooladzandi, Davini, and Mirzasoaleiman 2022).

We select the Core-Template using a greedy differentiable farthest point sampling (FPS) process with a learned metric parameterized by the quality of the face images. As computing the farthest point requires an Argmax operation, which is not differentiable, we suggest a novel differential FPS by replacing the Argmax operation with sampling from the differentiable Gumbel-Softmax (Jang, Gu, and Poole 2016) distribution of quality-aware distances. Unlike conventional FPS which starts with a point selected at random, our selection process starts with highest quality feature, hence the selection is permutation invariant. An example of a face template and the selected Core-Template is depicted in Fig. 1. As far as we know, this is the first differential coreset approach that can be fully optimized by backpropagation. Importantly, our differentiable Core-Template selection reduces the unbounded template size N to a fixed size K , facilitating the overall linear algorithm complexity.

To enhance the Core-Template representation, we compute self-attention amongst the Core-Template features. To enrich the representation with information from the full template missing from the Core-Template, we compute cross-attention between the Core-Template and the full template, using the Core-Template features as queries and the full template as Key-Values. By using the Core-Template features as queries for self-attention and cross-attention, we compute attention in linear time, unlike the typical quadratic time complexity in other works modeling intra-template relations, such as in the well-written paper by (Kim et al. 2022).

As a final step, we sum and normalize the enriched Core-Template features to get the final feature representation for the template. The final template representative feature is used to compute the vanilla FR adaptive margin loss (Kim,

Jain, and Liu 2022) to optimize the model, which is simple and effective. The Core-Template selection process and attention computation are linear in the template size, facilitating an overall linear time complexity of FaceCoresetNet. The proposed scheme is shown to achieve SOTA accuracy on the IJB-B (Whitelam et al. 2017), IJB-C (Maze et al. 2018) FSR benchmarks.

In particular, we propose the following contributions:

- We propose a novel differential, permutation invariant, coreset selection for SFR, with a differentiable FPS strategy based on a quality-aware metric to balance quality and diversity. The coreset selection and metric parameters are part of the trained model and optimized end-to-end. As far as we know, this is the first differential coreset selection method with potential value to other fields.
- We model intra-template relations using the selected coreset as queries in self and cross-attention architecture to extract information from the whole set.
- As far as we know, our algorithm is the first to compute intra-template relations in linear time.
- We use a vanilla FR AdaFace loss with no additional loss terms for a simple design.
- We establish a new SoTA for face verification on IJB-B and IJB-C datasets

Related Work

A simple method to combine features involves computing the mean of a collection of features (Chen et al. 2015; Parkhi, Vedaldi, and Zisserman 2015). A limitation of this elementary technique is that inferior-quality images can dominate the combined descriptor, diminishing its discriminatory capacity. During training with margin-based softmax, superior-quality faces indirectly produce features with greater magnitude (Meng et al. 2021; Kim, Jain, and Liu 2022). By summing the features, high-quality images are inherently emphasized without needing explicit image quality evaluation. Both CFAN (Deng et al. 2019) and QAN (Liu, Yan, and Ouyang 2017) employ a self-attention mechanism to learn image weighting through an explicit quality assessment (Gong et al. 2019; Liu, Yan, and Ouyang 2017). Nevertheless, these approaches share a common drawback: they do not consider the relationships within the set during the weight computation stage.

To model intra-set relationships, the nonlocal neural network and RSA (Liu et al. 2019; Wang et al. 2018) utilize intermediate feature maps U_i of size $C_m \times H \times W$, during aggregation. This approach leverages the rich and complementary information feature maps provide and considers spatial relationships to refine them. However, the main limitation is the intensive computation required for attention calculation. Specifically, for a set of N feature maps, an attention module requires an affinity map of size $(N \times H \times W)^2$. CAFace, cluster and Aggregate network (Kim et al. 2022) generates an affinity map of size N^2 , resulting in improved computation efficiency for attention calculation. Using our differential selection of coresets, this work reduces the feature fusion complexity to $O(N)$.

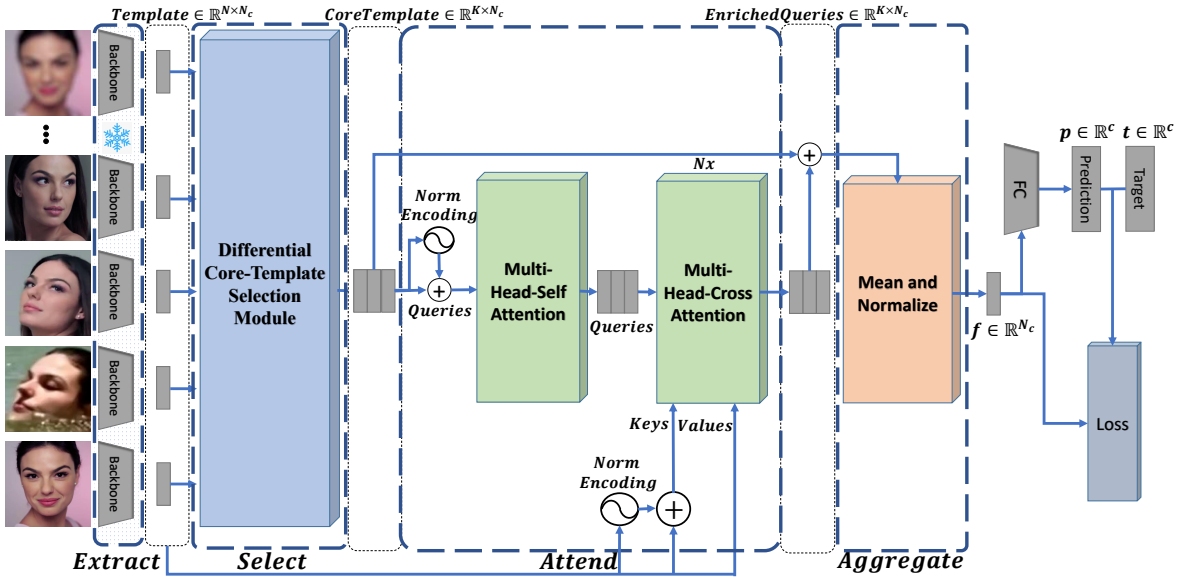


Figure 2: FaceCoresetNet Architecture

Set burstiness suppression network (Wang, Zhao, and Wu 2022), tackles the face burstiness, a problem mostly overlooked by FSR research, by quantizing the face set feature Gram matrix, which has a size of N^2 . In the main FSR datasets IJB-B, IJB-C, each image in each template is labeled by its media source (either a video clip or a singular image). MagFace and AdaFace (Kim, Jain, and Liu 2022; Meng et al. 2021) workaround the face burstiness problem by using these labels to fuse each media source separately and then sum the per-media features to produce the final fused feature. In this method, each media source contributes equally to the final descriptor, regardless of size, mitigating the burstiness problem. In our work, we solve face burstiness in linear time without resorting to media labels, which may not be available, and still achieve superior performance.

DAC (Liu et al. 2018) proposes an RL-based quality estimator, while MARN (Gong, Shi, and Jain 2019) is an RNN-based quality estimator. However, these methods are not agnostic to input order, making them unsuitable for modeling orderless sets. In contrast, our method is permutation invariant, and well suited for set modeling.

A long-standing practice in fundamental approaches such as kNN and kMeans (Har-Peled and Kushal 2005) or mixture models (Feldman, Faulkner, and Krause 2011) is the use of subsets of the core to approximate the structure of an available set. These subsets enable the finding of approximate solutions with considerably reduced costs (Agarwal et al. 2005). Recently, coreset-based methods have made their way into Deep Learning techniques, such as network pruning (Mussay et al. 2021), active learning (Kim and Shin 2022), and increasing effective data coverage of minibatches for improved GAN training (Sinha et al. 2020), or representation learning (Roth et al. 2020). The latter two have achieved success through a greedy coreset selection mechanism.

As far as we know, all previous coreset selection methods are non-differentiable. In contrast, in our proposed method, the coreset selection is integrated into our PyTorch model and trained end-to-end. Quality-based distance metric parameters are optimized during training, allowing an optimal balance between quality and diversity in the selected coreset.

To implement our differentiable Core-Template selection, we use the Gumbel-Softmax (Jang, Gu, and Poole 2016) technique, which enables the smoothing of discrete categorical distributions, making them suitable for backpropagation. It employs the Gumbel-Max reparametrization trick, which efficiently draws samples from a Categorical distribution. Let z be the categorical variable with class probabilities $\pi_1, \pi_2 \dots \pi_k$. The trick asserts that sampling from the categorical distribution is equivalent to sampling from the following expression:

$$z = \text{one_hot}(\underset{i}{\operatorname{argmax}}[g_i + \log \pi_i])$$

Here, $\log \pi_i$ denotes the logits of the categorical distribution, and g_i follows a Gumbel distribution with parameters (0,1). To avoid non-differentiability introduced by the Argmax operation, it is replaced with a Softmax function:

$$y_i = \frac{\exp((\log(\pi_i) + g_i)/\tau)}{\sum_{j=1}^k \exp((\log(\pi_j) + g_j)/\tau)} \text{ for } i = 1, \dots, k.$$

The temperature parameter (τ) controls the spreading of the distribution. As τ approaches zero, the Gumbel-Softmax method converges to the discrete categorical distribution. During training, a τ value of 1 is typically used, while during inference, τ approaches zero ($\tau \rightarrow 0^+$). (Yang et al. 2019) employ the Gumbel-Softmax for a differentiable sampling of point clouds, facilitating tasks such as point cloud segmentation and classification.

Multi-head self-attention (MSA) (Vaswani et al. 2017) is a commonly used function representing intra-set and inter-set relationships through an affinity map. It is a critical

component of transformer architectures, which have surpassed CNNs in various visual tasks (Dosovitskiy et al. 2021; Zhang et al. 2022; Zheng et al. 2021). In principle, each MSA output value is a weighted average of other values controlled by the queries and keys affinity. The MSA time complexity is $O(N^2)$; hence computing intra-set attention on a large template is prohibitive. To reduce the computation time, we compute intra-set attention on a fixed-size coreset, and cross-attention between the coreset and the template in $O(N)$ time.

Method

Let $T = \{x_1, x_2, \dots, x_N\}$ be a template of N facial images of the same identity. The goal is to produce a discriminative single compact feature vector f from T . To compute f , we employ an *Extract, Select, Attend and Aggregate* paradigm. To *Extract* (fig. 2 left) compact feature f_i from each image x_i , we use a pre-trained frozen single image FR model $E : x_i \rightarrow f_i$ following (Meng et al. 2021; Kim, Jain, and Liu 2022). The set of features produced is $F = \{f_1, f_2, \dots, f_N\}$ or in a matrix notation $F \in \mathbb{R}^{N \times N_c}$ where N_c is the number of channels in the compact feature. During training, we create a batch of B templates of B different individuals. Each template of size N where $N_{min} \leq N \leq N_{max}$, i.e. $F \in \mathbb{R}^{B \times N \times N_c}$. The template size N is randomly selected for each mini-batch.

Core Template Selection

To reduce the template feature size F from unbounded size N , we *Select* (Fig. 2 blue block) a core template denoted by CT of fixed size K . Our core template selection process strives to select high-quality images with diverse appearances to address the burstiness problem. As established in (Kim, Jain, and Liu 2022), $\|f_i\|$ is a proxy for x_i 's image quality. Our core template selection starts by selecting the feature f_i with the largest norm in F , the template's highest quality feature. Following the farthest point sampling (Eldar et al. 1997) paradigm, in each selection iteration, we choose the farthest feature in $F \setminus CT$ from the set of features already selected in CT . The farthest point in our setting should correlate to a high-quality image with a diverse appearance relative to the features already selected for the core template. Cosine distance between two features f_i and f_j is defined as:

$$d_c(f_i, f_j) := 1 - \frac{f_i \cdot f_j}{\|f_i\| \|f_j\|} \quad (1)$$

Large cosine distance between features correlates with large appearance variations (face pose, expression, age, etc.) between the corresponding images. Selecting the farthest feature using this metric encourages core template appearance diversification, but ignores image quality, resulting in the selection of diverse but low-quality images. To account for image quality, we scale the cosine distance by the feature norm exponentiated by a learned parameter γ to balance between quality and diversity:

$$d_q(f_i, f_j; \gamma) := d_c(f_i, f_j) \|f_j\|^\gamma \quad (2)$$

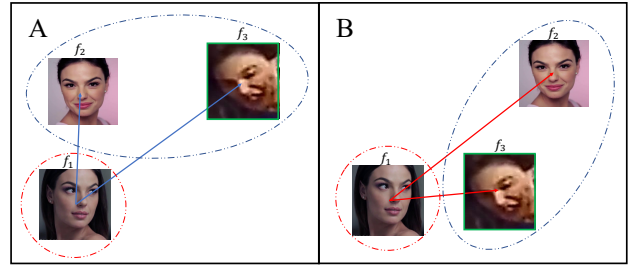


Figure 3: Cosine distance (A) vs. quality-aware distance (B) for core template selection. The distance from f_1 to f_2 increases in B relative to A, as f_2 has high quality. The distance from f_1 to f_3 is decreased in B, as f_3 has low quality.



Figure 4: The core template selection is influenced by the value of γ . In this example, the template size $N = 5$ and $K = 3$. In the top row, when γ is large, the selection process prioritizes image quality over diversity. As a result, three similar high-quality images are chosen, marked with red rectangles. Conversely, in the bottom row, with $\gamma = 0$, diversity precedes image quality, leading to a selection that emphasizes a diverse range of images.

Here, $f_i \in CT$ and $f_j \in F \setminus CT$ and d_q is a non-symmetric distance function modulated by image x_j 's quality. γ is a trained parameter that reflects the optimal importance of the image quality. If γ is 0, there is no weight to quality and d_q is reduced to d_c . If $\gamma \gg 1$, the feature quality dominates, and we have $d_q(f_i, f_j; \gamma) \approx \|f_j\|^\gamma$ as illustrated in Fig. 4. The effect of using our quality aware distance function vs. regular cosine distance is illustrated in Fig. 3. In Fig. 3A, we have the template $F = \{f_1, f_2, f_3\}$ and f_1 already in the Core Template ($CT = \{f_1\}$), marked by the red dashed ellipse, and we want to choose the next feature for the core template out of $F \setminus \{f_1\} = \{f_2, f_3\}$ (marked by the blue dashed ellipse). Using the regular cosine distance, the feature f_3 of low-quality image 3 is farther from f_1 than high-quality image 2 due to its larger pose difference. As a result, the low quality f_3 feature is selected for CT ($CT = \{f_1, f_3\}$), reducing the discriminative power of CT . In Fig. 3B, using our quality aware metric Eq. 2, the distance $d_c(f_1, f_2)$ is increased by f_2 's larger quality measure, and the distance $d_c(f_1, f_3)$ is reduced by f_3 lower quality, reversing the distance order and selecting the high quality f_2 for the core template ($CT = \{f_1, f_2\}$)

The Core-Template selection algorithm extends the FPS algorithm with important modifications. Instead of relying

Listing 1: Differential Core-Template Selection

```

1  def CT_selection(F, K, F_norms, gamma):
2      """
3          F: normalized features data, [B,
4              N, C]
5          K: The intended coreTemplate (
6              CT) size (scalar)
7          F_norms: F norms, [B, N]
8          gamma: learned quality vs.
9              diversity balance parameter (
10                 scalar)
11
12         Return:
13             CoreTemplate CT with size K
14         """
15         tau = 1.0 if train else 1e-10
16         l_hot_max = torch.gumbel_softmax(
17             F_norms, tau=tau, hard=True)
18         CT = F @ l_hot_max
19         d_CT_to_F = quality_aware_dist(CT,
20             F_norms, F)
21         for i in range(K-1):
22             l_hot_max = torch.gumbel_softmax(
23                 (d_CT_to_F, tau=tau, hard=
24                     True)
25             new_f = F @ l_hot_max
26             d_new_f_to_CT =
27                 quality_aware_dist(new_f,
28                     norms, F, gamma)
29             d_CT_to_F = torch.min(d_CT_to_F,
30                 d_new_f_to_CT)
31             CT = torch.cat([CT, new_f])
32         return CT
33
34 def quality_aware_dist(candidate_point,
35     F_norms, F, gamma):
36     inner_product = torch.bmm(F,
37         candidate_point)
38     cosine_dist = (1 - inner_product)
39     quality_dist = torch.pow(norms,
40         gamma) * cosine_dist
41     return quality_dist

```

on the non-differentiable Argmax operation to choose the feature in T farthest from CT , we adopt a differentiable approach by sampling from the Gumbel-Softmax distribution of distances (Jang, Gu, and Poole 2016). This enables us to incorporate differentiability and leverage backpropagation during training. Additionally, we introduce a quality-based distance function that incorporates a learned parameter γ , allowing the algorithm to achieve an optimal balance between quality and diversity in the selected Core-Template features. Furthermore, unlike the conventional FPS method that randomly selects the first feature, we prioritize the feature with the highest norm (highest quality) to be inserted first into the Core-Template CT . This selection criterion ensures permutation invariance and enhances the algorithm’s overall performance. The pseudocode for our core template selection algorithm is in Listing 1.

The algorithm starts by sampling from the Gumbel-Softmax distribution of the template’s norms F (line 11). The sampling operator returns a 1-hot vector that is used

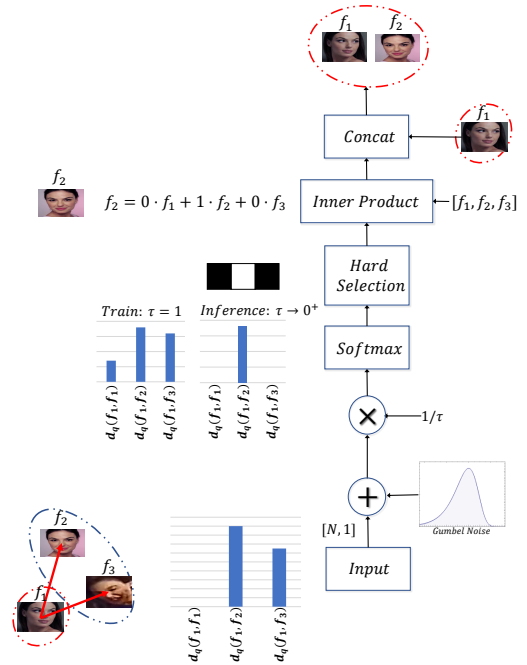


Figure 5: Illustration of a single iteration of the differential core-template selection (lines 14 to 19) in Listing. 1

in line 12 to insert the highest quality feature to the Core-Template CT . In line 13, we calculate the distance from each feature in F to the newly selected feature in CT . In each iteration of the loop starting at line 14 we add one new member to CT that is farthest from it. We use the quality aware distance function in line 22 for all the distance calculations. An example of one iteration in this process is depicted in Figure 5. Consider the feature set $F = \{f_1, f_2, f_3\}$ and the initial Core-Template $CT = \{f_1\}$. The quality aware distances from F to CT , serve as input to the selection process. In the next stage, Gumbel noise is added to each distance value. These values are then multiplied by the temperature parameter τ , and the softmax operation is applied to obtain the final Gumbel-Softmax distribution. A mask is generated to identify the maximal value in the distribution, which is then utilized to determine the next feature to add to CT .

Importantly, our proposed core template selection process time complexity is $O(NK)$ where N is the template size, and K is the core template size with $K \ll N$. Assuming K is a small constant (we found $K = 3$ optimal in our experiments), we have a selection process with linear complexity. In addition, the core template selection process is permutation invariant, as it depends solely on the geometrical attributes of each feature in F and not on their order.

Attend

Our core template selection process performs an important data reduction from unbounded template size N to a small fixed-size core template. This data reduction enables the application of attention mechanisms like multi-head-self-

attention (Vaswani et al. 2017), which have quadratic time complexity and are prohibitively slow for large templates.

We use Multihead Self Attention (MSA) as an encoder to enhance the selected core template representation power and exchange information among its features (Fig. 2).

Inspired by recent developments in detection transformers such as Efficient-DETR (Yao et al. 2021), which demonstrate the superiority of leveraging image features as queries instead of relying on learned queries for decoding as in the original transformer (Vaswani et al. 2017) and DETR (Carion et al. 2020), we leverage the core template features as transformer decoder queries that interact with key values taken from the full template features in F . This allows the core template features to extract additional information from the full template that is missing from the core template.

Norm encoding. It appeals to our intuition that explicit image quality data can aid information extraction in multi-head attention settings, facilitating the suppression of low-quality features and highlighting high-quality features. To incorporate the quality information, we encode the feature norms as a quality proxy, using the standard Sinusoidal conversion established in (Vaswani et al. 2017), and add them to the original features. Our experiments validate the positive impact of Norm encoding on the model’s overall accuracy.

Importantly, our attention process interpolates between template features using self and cross-attention realized with MHA. We avoid the Feed-Forward block in the transformer encoder and decoder blocks, which extrapolates the template features and deteriorates performance.

The complexity of the encoder is $O(K^2)$, which is constant. The complexity of the decoder is $O(NK)$. The total complexity of the attention block is $O(N)$.

Aggregate. To represent the entire template with one feature f , we combine the K features in the enhanced core template by adding them up and normalizing the result to have a unit length.

Training and Loss

After the aggregation phase, we get a single vector f representing the whole template. From this point on, the problem is reduced from SFR to FR, facilitating the use of SoTA FR loss functions. In this work, we choose AdaFace (Kim, Jain, and Liu 2022) loss for its superior accuracy. AdaFace computes a class margin adaptively based on the image quality of the face approximated by the face feature norm. In our case, the margin is computed by the *template* quality instead of the image (See Fig. 2 right-hand side). Using a simple margin-based loss for template feature fusion facilitates a simplistic implementation compared with elaborate loss terms found in (Kim, Jain, and Liu 2022) and others.

Large modern FR datasets like WebFace (Zhu et al. 2021) contain millions of identities but only a limited number of images per identity (WebFace4M contains about 20 images per identity on average) with no video clips. On the other hand, SFR benchmarks like IJB-B, IJB-C contain rich templates with a diverse combination of still images and video clips for each identity. To bridge this gap, we online augment the training set and simulate IJB-B, IJB-C style templates. To compensate for the lack of video clips, we simulate a

#	Differential Core-Template Selection	Self Att	Cross Att	Norm Encoding	IJB-B:TAR@ FAR=1e-6
1	×	×	×	×	38.45
2	✓	×	×	×	48.01
3	✓	✓	×	×	49.24
4	✓	✓	✓	×	51.71
5	✓	✓	✓	✓	52.56

Table 1: FaceCoresetNet Ablations

video clip by randomly sampling an image and applying various geometrical transformations (translations, rotations, etc.) on it to create multiple “video frames”. Each training template comprises a random collection of still photos and simulated video clips with random addition of photometric noise to each image.

Experiments

To allow for a valid comparison with the current SOTA, we strictly follow the experiments protocol of CAFace (Kim et al. 2022) (NeurIPS’22). Our training dataset is WebFace4M (Zhu et al. 2021), which contains 4.2 million facial images from 205,990 identities. We use CAFaces’s pre-trained face recognition model, E , an IResnet-101, trained with AdaFace loss on the entire WebFace4M dataset. To train the FaceCoresetNet, we use CAFace’s subset of 813,482 images from 10,000 identities. As VGG-2 (Cao et al. 2018) and MS1MV2 (Guo et al. 2016; Deng et al. 2019) creators have withheld them due to privacy and other issues, we do not use them in our experiments. We train FaceCoresetNet for 2 epochs till convergence, compared with CAFace’s 10 epochs training schedule.

We test on IJB-B (Whitelam et al. 2017) and IJB-C (Maze et al. 2018) datasets as they are intended for template-based face recognition. IJB-B comprises 1845 individuals and includes 11,754 images, 55,025 frames, and 7,011 videos. Each template in the dataset consists of a diverse collection of still images and video frames from various sources. These images and videos were gathered from the Internet and are considered unconstrained, as they exhibit considerable variation in pose, lighting conditions, and image quality. Following CAFace, we focus our experiments on face verification, where IJB-B offers 10,270 genuine template pair comparison and 8,000,000 impostor comparisons, facilitating valid measurements of very small FAR values on the TAR@FAR ROC curve. IJB-C extends IJB-B by adding 1,661 new subjects, with increased emphasis on occlusion and diversity.

Ablation and Analysis

To demonstrate the effectiveness of each component in FaceCoresetNet, we conduct a series of experiments by ablating individual components. Initially, we turned off all components, resulting in a configuration where the template feature was simply the average feature computed from all the templates’ features. The obtained TAR@FPR=1e-6 for the IJB-B dataset was 38.45 (Table 1). Subsequently, we enabled the differential core-template selection. We computed

Method	Complexity (N : Template Size)	FAR=1e-4	IJB-B 1:1 Verification TAR			IJB-C 1:1 Verification TAR	
			FAR=1e-5	FAR=1e-6	FAR=1e-4	FAR=1e-5	FAR=1e-6
AdaFace (Kim, Jain, and Liu 2022)	NA	94.84	90.86	38.45	96.42	94.47	87.90
RSA (Liu et al. 2019)	$O(N^2)$	95.00	91.22	-	96.49	94.58	-
SD + VBA (Wang, Zhao, and Wu 2022)	$O(N^2)$	95.38	90.89	48.96	96.65	94.85	90.02
CAFace (Kim et al. 2022)	$O(N^2)$	95.78	92.78	47.21	97.30	95.96	90.56
FaceCoresetNet	$O(N)$	95.28	91.44	52.56	96.96	95.21	91.12

Table 2: A performance comparison of recent methods on IJB-B (Whitelam et al. 2017) and IJB-C (Maze et al. 2018) datasets. FaceCoresetNet achieves the best results on TPR@FPR=1e-6 for both IJB-B and IJB-C while being the most efficient.

γ policy	IJB-C: TAR@FAR=1e-6
Fixed 0 (Diversity priority)	90.84
Fixed 10 (Quality priority)	89.69
Trained	91.12

Table 3: Training γ for optimal balance between quality and diversity yields best results, compared with fixed γ values

the final template descriptor by averaging the pooled features within the core template. Surprisingly, by effectively selecting 3 features, the accuracy increased significantly to 48.01. We hypothesize that our core template’s high quality and diversity reduce the amount of noise present in the original template, which may explain the observed improvement in accuracy. Furthermore, we integrated self-attention into the core-template features, boosting accuracy to 49.24. Additionally, by employing cross-attention between the core and full templates, the accuracy increased to 51.71. Finally, with the incorporation of Norm encoding, the accuracy reached 52.56.

Effect of Core-Template size. To determine the optimal core-template size, we searched a range of sizes from 1 to 6. The outcomes of this search are summarized in Table 1 of the supplementary materials. Notably, the ideal coresets size was identified as 3, and this optimal value remained constant throughout all our subsequent experiments.

γ computation policy. Our model optimizes the γ value in the quality-aware distance. To showcase the effectiveness of this optimization, we conducted 3 experiments. One setting γ to the constant 10 value, prioritizing quality over diversity. Conversely, we set γ to 0, which reduces the quality aware distance to the vanilla cosine distance, prioritizing diversity over quality. In the last experiment, we optimize γ end-to-end. The results in Table 3 demonstrates the effectiveness of our approach.

Comparison with SoTA methods. In Table. 2, we compare our results with the state-of-the-art (SOTA) in face verification, focusing on the IJB-B and IJB-C datasets. Our method achieves the highest accuracy in the low false acceptance rate regime, TAR@FAR=10e-6, for both IJB-B and IJB-C, outperforming others significantly. Additionally, we obtain the second-best results for TAR@FAR=10e-5 while maintaining the advantage of utilizing the lowest **linear** time compute.




Figure 6: Computation complexity comparison with SoTA. FaceCoresetNet is linear in the template size, while CAface is quadratic.

Computation Complexity. To validate our algorithm’s claimed linear time complexity, we employ (fvcore 2023) to calculate the FLOPs in our FaceCoresetNet and compare it against CAface (Kim et al. 2022) using the authors’ code. The FLOPs measurements for different template sizes are illustrated in Fig. 6. The results demonstrate that our algorithm indeed operates in linear time, while CAface exhibits quadratic time complexity.

Conclusion

In this work, we employ a method to fuse face template features by selecting a small, fixed-size core-template that balances quality and diversity. This selection uses a differentiable farthest-point sampling with a quality-aware distance metric, realized by sampling from the Gumbel-Softmax distribution of quality-aware distances.

The selected core template features are further enhanced using MHA self-attention. These enhanced features are then used as queries for MHA cross-attention between the core template and the entire template, enabling the extraction of additional information that may be missing from the core template.

Our proposed model achieves state-of-the-art accuracy on IJB-B and IJB-C datasets, all while maintaining reduced linear time complexity and simple design. FaceCoresetNet is showcased within the context of face set recognition, yet its applicability may extend to other domains requiring efficient feature fusion based on a differentiable selection with a learned metric. We hope to see an extension of this work to other domains in future research.

References

- Agarwal, P. K.; Har-Peled, S.; Varadarajan, K. R.; et al. 2005. Geometric approximation via coresets. *Combinatorial and computational geometry*, 52(1): 1–30.
- Campbell, T.; and Broderick, T. 2018. Bayesian coreset construction via greedy iterative geodesic ascent. In *Proceedings of the International Conference on Machine Learning (ICML)*, 698–706. PMLR.
- Cao, Q.; Shen, L.; Xie, W.; Parkhi, O. M.; and Zisserman, A. 2018. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition (FGR)*, 67–74. IEEE.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 213–229. Springer.
- Chen, J.-C.; Ranjan, R.; Kumar, A.; Chen, C.-H.; Patel, V. M.; and Chellappa, R. 2015. An end-to-end system for unconstrained face verification with deep convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, 118–126.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. ArcFace: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4690–4699.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Eldar, Y.; Lindenbaum, M.; Porat, M.; and Zeevi, Y. Y. 1997. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9): 1305–1315.
- Feldman, D.; Faulkner, M.; and Krause, A. 2011. Scalable Training of Mixture Models via Coresets. In Shawe-Taylor, J.; Zemel, R.; Bartlett, P.; Pereira, F.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems (NIPS)*, volume 24. Curran Associates, Inc.
- fvcore. 2023. fvcore. <https://github.com/facebookresearch/fvcore>.
- Gong, S.; Shi, Y.; and Jain, A. 2019. Low quality video face recognition: Multi-mode aggregation recurrent network (MARN). In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 0–0.
- Gong, S.; Shi, Y.; Kalka, N. D.; and Jain, A. K. 2019. Video face recognition: Component-wise feature aggregation network (c-fan). In *International Conference on Biometrics (ICB)*, 1–8. IEEE.
- Graves, A.; and Graves, A. 2012. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, 37–45.
- Guo, Y.; Zhang, L.; Hu, Y.; He, X.; and Gao, J. 2016. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 87–102. Springer.
- Har-Peled, S.; and Kushal, A. 2005. Smaller coresets for k-median and k-means clustering. In *Proceedings of the twenty-first annual symposium on Computational geometry*, 126–134.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Kalka, N. D.; Maze, B.; Duncan, J. A.; O’Connor, K.; Elliott, S.; Hebert, K.; Bryan, J.; and Jain, A. K. 2018. Ijb-s: Iarpa janus surveillance video benchmark. In *2018 IEEE 9th international conference on biometrics theory, applications and systems (BTAS)*, 1–9. IEEE.
- Kim, M.; Jain, A. K.; and Liu, X. 2022. AdaFace: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18750–18759.
- Kim, M.; Liu, F.; Jain, A.; and Liu, X. 2022. Cluster and Aggregate: Face Recognition with Large Probe Set. In *Advances in Neural Information Processing Systems (NIPS)*.
- Kim, Y.; and Shin, B. 2022. In Defense of Core-set: A Density-aware Core-set Selection for Active Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 804–812.
- Liu, X.; Guo, Z.; Li, S.; Kong, L.; Jia, P.; You, J.; and Kumar, B. 2019. Permutation-invariant feature restructuring for correlation-aware image set-based recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 4986–4996.
- Liu, X.; Kumar, B.; Yang, C.; Tang, Q.; and You, J. 2018. Dependency-aware attention control for unconstrained face recognition with image sets. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 548–565.
- Liu, Y.; Yan, J.; and Ouyang, W. 2017. Quality aware network for set to set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5790–5799.
- Maze, B.; Adams, J.; Duncan, J. A.; Kalka, N.; Miller, T.; Otto, C.; Jain, A. K.; Niggel, W. T.; Anderson, J.; Cheney, J.; and Grother, P. 2018. IARPA Janus Benchmark - C: Face Dataset and Protocol. In *International Conference on Biometrics (ICB)*, 158–165.
- Meng, Q.; Zhao, S.; Huang, Z.; and Zhou, F. 2021. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14225–14234.
- Mirzasoleiman, B.; Bilmes, J.; and Leskovec, J. 2020. Coresets for data-efficient training of machine learning models. In *International Conference on Machine Learning*, 6950–6960. PMLR.
- Mussay, B.; Feldman, D.; Zhou, S.; Braverman, V.; and Osadchy, M. 2021. Data-independent structured pruning of neural networks via coresets. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12): 7829–7841.
- Parkhi, O. M.; Vedaldi, A.; and Zisserman, A. 2015. Deep face recognition.

- Pooladzandi, O.; Davini, D.; and Mirzasoleiman, B. 2022. Adaptive second order coresets for data-efficient machine learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 17848–17869. PMLR.
- Roth, K.; Milbich, T.; Sinha, S.; Gupta, P.; Ommer, B.; and Cohen, J. P. 2020. Revisiting training strategies and generalization performance in deep metric learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 8242–8252. PMLR.
- Sinha, S.; Zhang, H.; Goyal, A.; Bengio, Y.; Larochelle, H.; and Odena, A. 2020. Small-gan: Speeding up gan training using core-sets. In *Proceedings of the International Conference on Machine Learning (ICML)*, 9005–9015. PMLR.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems (NIPS)*, 30.
- Wang, J.; Zhao, Z.; and Wu, F. 2022. Set-Based Face Recognition Beyond Disentanglement: Burstiness Suppression With Variance Vocabulary. In *Proceedings of the 30th ACM International Conference on Multimedia*, 6125–6135.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7794–7803.
- Whitelam, C.; Taborsky, E.; Blanton, A.; Maze, B.; Adams, J.; Miller, T.; Kalka, N.; Jain, A. K.; Duncan, J. A.; Allen, K.; Cheney, J.; and Grother, P. 2017. IARPA Janus Benchmark-B Face Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 592–600.
- Xie, W.; and Zisserman, A. 2018. Multicolumn networks for face recognition. *arXiv preprint arXiv:1807.09192*.
- Yang, J.; Zhang, Q.; Ni, B.; Li, L.; Liu, J.; Zhou, M.; and Tian, Q. 2019. Modeling point clouds with self-attention and gumbel subset sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3323–3332.
- Yao, Z.; Ai, J.; Li, B.; and Zhang, C. 2021. Efficient detr: improving end-to-end object detector with dense prior. *arXiv preprint arXiv:2104.01318*.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H. 2022. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Zheng, S.; Lu, J.; Zhao, H.; Zhu, X.; Luo, Z.; Wang, Y.; Fu, Y.; Feng, J.; Xiang, T.; Torr, P. H.; et al. 2021. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6881–6890.
- Zhu, Z.; Huang, G.; Deng, J.; Ye, Y.; Huang, J.; Chen, X.; Zhu, J.; Yang, T.; Lu, J.; Du, D.; et al. 2021. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10492–10502.