

HP-Capsule: Unsupervised Face Part Discovery by Hierarchical Parsing Capsule Network

Chang Yu^{1,2}, Xiangyu Zhu^{1,2}, Xiaomei Zhang^{1,2}, Zidu Wang^{1,2}, Zhaoxiang Zhang^{1,2,3}, Zhen Lei^{1,2,3*}

¹NLPR, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation, Chinese Academy of Sciences

{chang.yu, xiangyu.zhu, zlei}@nlpr.ia.ac.cn

{zhangxiaomei2016, wangzidu2022, zhaoxiang.zhang}@ia.ac.cn

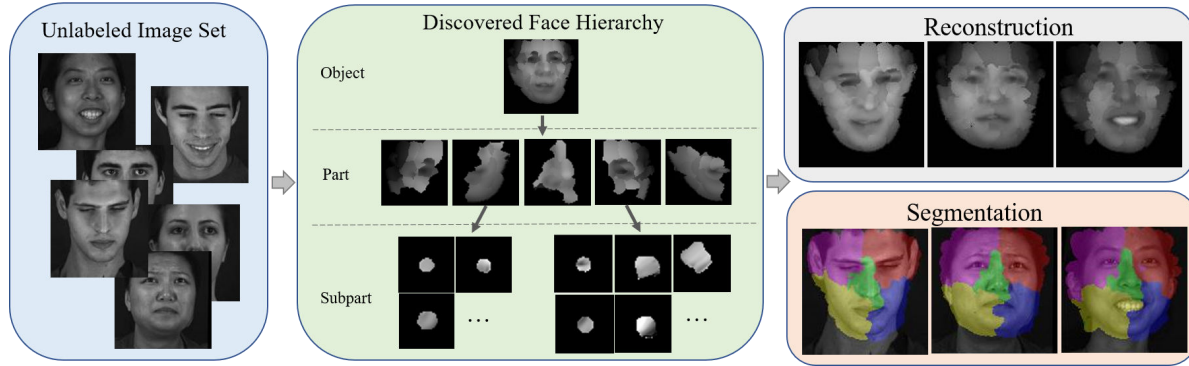


Figure 1. A brief review of the Hierarchical Parsing Capsule Network (HP-Capsule). Given a large scale of unlabeled images (left), HP-Capsule can automatically discover the hierarchical face parts (middle) and give unsupervised face segmentation results as by-products (right).

Abstract

Capsule networks are designed to present the objects by a set of parts and their relationships, which provide an insight into the procedure of visual perception. Although recent works have shown the success of capsule networks on simple objects like digits, the human faces with homologous structures, which are suitable for capsules to describe, have not been explored. In this paper, we propose a Hierarchical Parsing Capsule Network (HP-Capsule) for unsupervised face subpart-part discovery. When browsing large-scale face images without labels, the network first encodes the frequently observed patterns with a set of explainable subpart capsules. Then, the subpart capsules are assembled into part-level capsules through a Transformer-based Parsing Module (TPM) to learn the compositional relations between them. During training,

as the face hierarchy is progressively built and refined, the part capsules adaptively encode the face parts with semantic consistency. HP-Capsule extends the application of capsule networks from digits to human faces and takes a step forward to show how the neural networks understand homologous objects without human intervention. Besides, HP-Capsule gives unsupervised face segmentation results by the covered regions of part capsules, enabling qualitative and quantitative evaluation. Experiments on BP4D and Multi-PIE datasets show the effectiveness of our method.

1. Introduction

Psychological studies [12, 22, 27] reveal that: The recognition procedure is often assigned with hierarchical structural descriptions by parsing the shapes into components and organizing them with their spatial relationships. This statement is consistent with the improvements of various

*Corresponding author.




Figure 2. The parts discovered by SCAE [17] on digits and faces. Even though SCAE can capture parts on simple digits, it tends to fail on more complicated objects like human faces, learning holistic templates as parts.

perceptual tasks which incorporate part-level information for better embeddings [3, 24, 37]. However, most of the existing methods implement the parsing in a predefined way [18, 20, 40, 43], where the definitions of parts are given by humans. Such handcrafted parsing can not reflect how neural networks understand objects.

To explore the visual perception mechanism of neural networks, an intriguing way is to learn the parsing directly from the data. The network should discover the visual part concepts with as little human intervention as possible and keep the semantic consistency across different samples, as shown in Figure 1. This unsupervised learning task is still a challenging problem, as the semantic parts are difficult to be described mathematically.

Capsule networks, which are designed to present objects by a set of parts and their relationships, are a feasible solution for this unsupervised parsing task. Among the capsule structures proposed in recent years [13, 14, 17, 25, 26], SCAE [17] is most suitable for this unsupervised face part discovery task since it presents objects with a part-whole hierarchy and defines the capsule as a set of explainable parameters including presences, poses, and visualizable templates. However, SCAE can only parse simple objects like handwritten digits. When extended to face images, SCAE fails to capture face parts and decomposes faces as a whole, generating holistic representations in the part capsules, as shown in Figure 2.

In this paper, we will improve the capsule network to handle human faces, which have homologous structures but diverse appearances. A Hierarchical Parsing Capsule Network (HP-Capsule) is proposed to discover hierarchical face parts and their relationships directly from the unlabeled image sets. Specifically, HP-Capsule understands faces with two sub-modules: a capsule-based autoencoder for subpart discovery and a Transformer-based Parsing Module (TPM) for hierarchy construction. During training, the frequently observed patterns are first captured and encoded by the capsule encoder with a set of explainable capsule parameters. A Visibility Activation Function (VAF) is

proposed to constrain the process of reconstructing the input image with capsule templates, so that the object is spatially decomposed into subparts rather than holistic representations. Then, the discovered subparts are regarded as visual words and sent to TPM to be aggregated to higher-level part capsules, where the subpart-to-part hierarchy is naturally built. Several constraints are incorporated to preserve the shape and appearance consistency so that the generated parts have more prominent semantics.

As a by-product, the covered regions of part capsules can be regarded as the segmentation maps, which can be used for the unsupervised face segmentation task, enabling the evaluation of our method. Compared with other unsupervised segmentation methods, HP-Capsule shows better semantic consistency and provides more interpretable descriptions about the discovered parts, including the visualizable templates and the statistics on presence and pose.

To summarize, the main contributions of this work are:

- This paper proposes a Hierarchical Parsing Capsule Network (HP-Capsule) for unsupervised face hierarchy discovery. HP-Capsule provides an insight into how the neural network understands homologous structures without human intervention.
- In the subpart discovery process, we propose a Visibility Activation Function (VAF), which enforces the network to concentrate on template regions with higher visibility, to ensure the objects are decomposed into localized subparts rather than holistic representations.
- A Transformer-based Parsing Module (TPM) is proposed to aggregate subparts into parts, constructing the subpart-part face hierarchy. The covered regions of part capsules can also be used for unsupervised face segmentation. Experiments on BP4D and Multi-PIE show the effectiveness of our method.

2. Related Work

Capsule Networks. Inspired by the sparse connection of human brains, capsule networks are designed to present the objects with a dynamic parse tree. Given an input image, the capsule network would automatically activate some of the capsules and then pass them to the next layers. When assigning the capsules with explicit meanings, the interpretability of the network would be improved. Recent works have shown the success of capsules on various tasks [28, 36, 38, 44] but most of them use capsules as the enhanced MLPs, while the interpretability, as well as the parsing characteristics of capsules, have not been well explored. For years, capsule networks have evolved many versions [13, 14, 17, 25, 26]. Among them, SCAE [17]

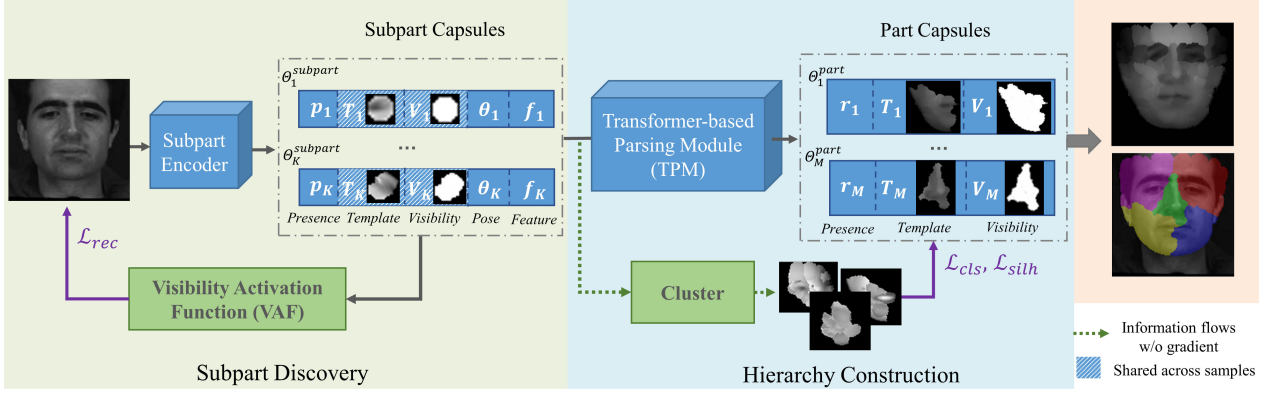


Figure 3. Overview of the Hierarchical Parsing Capsule Network (HP-Capsule). HP-Capsule understands faces with two sub-modules: the capsule autoencoder with Visibility Activation Function (VAF) for subpart discovery and the Transformer-based Parsing Module (TPM) for hierarchy construction. During training, the frequently observed patterns are firstly captured and encoded with a set of explainable capsule parameters, which are then passed through the VAF for subpart discovery. After that, the explored subpart capsules are sent into the TPM for higher-level part capsules. TPM is trained with the shape consistency loss and the pseudo subpart-part relationships generated by clustering. By automatically aggregating subparts into parts, the face hierarchy is naturally built.

is an intriguing structure that describes the objects with a set of visualizable templates through unsupervised learning. However, SCAE can only tackle simple objects like hand-written digits where the learned templates are stroke-like. Recently, Sabour *et al.* [26] extend the capsule network for complicated images like human bodies, but it needs optical flow as clues and can only show the mask of each part. In this paper, we extend SCAE to more complicated objects like human faces, discovering the face hierarchy directly from the unlabeled image sets.

Unsupervised Part Discovery. Unsupervised learning has witnessed impressive progress in recent years [2, 4, 10, 29, 31, 33–35]. Our work is related to the task of unsupervised part discovery, which aims to discover visual concepts from unlabeled images or videos. In the early years, Feng *et al.* [6] propose the local non-negative matrix factorization (LNMF) to learn part-based representations for human faces. By adding localized sparsity factorization, LNMF can learn part-based basis, but the performance is still constrained as the faces are presented in the linear space. Recent works [1, 8] show that the semantic parts have already been included in the CNN features. Inspired by this, Collins *et al.* [5] propose to use non-negative matrix factorization (NMF) on CNN activations to locate the semantic concepts on image sets. However, this approach needs to solve optimization during inference and lacks interpretability. Besides using activation maps, other methods [7, 26, 39] try to use motion clues in videos to discover parts as the regions with the same semantics always move together. Differently, our method does not need motion clues and the learned parts can be visualized directly to gain better explainability.

Unsupervised Face Segmentation. Recently, several

works have made positive progress on the unsupervised face segmentation task. Hung *et al.* [15] propose several loss functions to learn the face parts that are geometrically concentrated and robust to spatial transformations. Their method needs saliency maps to suppress the background features and Liu *et al.* [19] argue that this might make the method less reliable. As a result, they propose a concentration loss to separate the background and a squeeze-and-expand block for better shape representations. Gao *et al.* [7] propose a dual procedure by leveraging the motion clues embedded in videos to further improve the performance. Compared with these methods, our networks are more explainable and can obtain parts with better semantic consistency.

3. Methodology

Given a collection of images from the same category, our work aims to learn a model that can discover hierarchical parts and their relationships from a single image. In the following parts, we will introduce the overall framework in Section 3.1, the capsule autoencoder with Visibility Activation Function for subpart discovery in Section 3.2, and the Transformer-based Parsing Module for hierarchy construction in Section 3.3.

3.1. Overall Framework

In this paper, we propose a Hierarchical Parsing Capsule Network for unsupervised face hierarchy discovery, which includes a capsule autoencoder for subpart discovery and a parsing module for hierarchy construction. The overall framework is shown in Figure 3.

Given an input image I , the network first uses a capsule




Figure 4. The impact of the Visibility Activation Function (VAF). Compared with the holistic templates in (a), the subpart templates discovered in (b) are sparse and local-connected.

encoder to estimate K subpart capsules:

$$\Theta_1^s, \Theta_2^s, \dots, \Theta_K^s = \mathcal{E}_{enc}(\mathbf{I}), \quad (1)$$

where each subpart capsule Θ_k^s is a set of parameters with explainable physical meanings: $\Theta_k^s : \{p_k^s, \theta_k^s, T_k^s, V_k^s, f_k^s\}, k \in [1, K]$, including presence probability $p_k^s \in \mathbb{R}^{1 \times 1}$, pose $\theta_k^s \in \mathbb{R}^{1 \times 6}$, template $T_k^s \in \mathbb{R}^{H^s \times W^s}$, visibility map $V_k^s \in \mathbb{R}^{H^s \times W^s}$ and input-specific feature $f_k^s \in \mathbb{R}^{1 \times D}$. Particularly, T_k^s and V_k^s describe the shape and the visible region of the discovered subparts, which are shared among samples.

The discovered subpart capsules are regarded as visual words and are sent into a transformer-based parsing module to estimate higher-level part capsules, where the subpart-to-part hierarchy is naturally built:

$$\Theta_1^p, \Theta_2^p, \dots, \Theta_M^p = \mathcal{E}_{TPM}(\Theta_1^s, \Theta_2^s, \dots, \Theta_K^s), \quad (2)$$

where the m -th part capsule $\Theta_m^p : \{T_m^p, V_m^p, r_m^p\}$ includes template $T_m^p \in \mathbb{R}^{H^p \times W^p}$, visibility map $V_m^p \in \mathbb{R}^{H^p \times W^p}$ and relationship $r_m^p \in \mathbb{R}^{1 \times K}$ that describes the compositional relations between all subpart capsules. The T_m^p and V_m^p are aggregated from subparts according to the relationship r_m^p . **When the face hierarchy is constructed, the part capsules naturally encode the discovered face parts with semantics, whose covered regions can also be used for unsupervised face segmentation.**

3.2. Subpart Discovery by Visibility Activation Function

Parts can be roughly defined as sparse and local-connected regions that are semantically consistent across objects [8, 15]. The recent capsule network SCAE [17] has shown promising results on simple objects like digits, but tends to fail on more complicated objects like human faces. As shown in Figure 4a, SCAE captures holistic face representations as the templates in subpart capsules. To enable the learning of part-level templates, we propose

a Visibility Activation Function (VAF) for the capsule autoencoder.

During training, the network firstly uses a self-attention based capsule encoder [17] to estimate the parameters of subpart capsules $\Theta_k^s : \{p_k^s, \theta_k^s, T_k^s, V_k^s, f_k^s\}$. Secondly, the color component C_k^s is decoded from feature f_k^s for template refinement. Thirdly, template T_k^s and visibility map V_k^s are transformed by the pose parameter θ_k^s . Finally, the transformed \hat{T}_k^s and \hat{V}_k^s are passed through VAF to get the activated template \tilde{T}_k^s . This capsule encoding process can be formulated as:

$$\begin{aligned} p_k^s, \theta_k^s, T_k^s, V_k^s, f_k^s &= \mathcal{E}_{enc}(\mathbf{I}), \\ C_k^s &= \text{MLP}(f_k^s), \\ \hat{T}_k^s, \hat{V}_k^s &= \text{AffineTrans}(T_k^s, V_k^s, \theta_k^s), \\ \tilde{T}_k^s &= \text{VAF}(\hat{T}_k^s, \hat{V}_k^s). \end{aligned} \quad (3)$$

Considering the visibility map \hat{V}^s controls the silhouette of subparts, we employ VAF to suppress the learning of invisible regions, which is formulated as:

$$\tilde{T}_{k,i,j}^s = \begin{cases} \hat{T}_{k,i,j}^s, & \text{if } \hat{V}_{k,i,j}^s \geq \gamma \\ -1, & \text{otherwise,} \end{cases} \quad (4)$$

where i, j stands for the position of the pixel and γ is the hyper-parameter threshold for activation.

VAF encourages the network to reconstruct the input image only with the high-visibility template regions, so that the non-activated regions with low visibility, which are always represented as the blurred mean face, are suppressed without receiving back-propagated signals during training. By Eqn. 3, the templates of subpart capsules can be transformed into the image space. If the input image \mathbf{I} is modeled as a Gaussian Mixture of subpart capsules, the reconstruction loss is given by:

$$\mathcal{L}_{rec} = - \sum_{i,j} \sum_{k=1}^K p_k^s \hat{V}_{k,i,j}^s \mathcal{N}(I_{i,j} | C_{k,i,j}^s \cdot \tilde{T}_{k,i,j}^s; \sigma_k^2), \quad (5)$$

where σ_k^2 presents the variance of the Gaussian Mixture.

We also incorporate \mathcal{L}_{pres} to constrain the sparsity of the activated subpart capsules:

$$\mathcal{L}_{pres} = \frac{1}{B} \sum_{b=1}^B \left(\sum_{k=1}^K p_{k,b}^s - \tau \right)^2 + \frac{1}{K} \sum_{k=1}^K \left(\sum_{b=1}^B p_{k,b}^s - \frac{\tau B}{K} \right)^2, \quad (6)$$

where B stands for the batch size and τ is the hyper-parameter for the average number of activated subpart capsules. The network optimized by Eqn. 6 will try to activate the same number but different subpart capsules during the mini-batch.

Besides, we employ two losses to constrain the connectivity of the discovered subparts, including the \mathcal{L}_{cen} for

geometrical concentration [15] and \mathcal{L}_{std} for balancing the visible regions of different subpart templates:

$$\begin{aligned}\mathcal{L}_{cen} &= \sum_k \sum_{i,j} \|(i,j) - (W^s/2, H^s/2)\|_2 \cdot V_{k,i,j}^s, \\ \mathcal{L}_{std} &= std\left(\sum_{i,j} V_{k,i,j}^s\right),\end{aligned}\quad (7)$$

where W^s and H^s are the width and height of subpart templates, std stands for the standard deviation.

The loss function for subpart discovery is combined as:

$$\mathcal{L}_{subpart} = \lambda_{rec}\mathcal{L}_{rec} + \lambda_{pres}\mathcal{L}_{pres} + \lambda_{cen}\mathcal{L}_{cen} + \lambda_{std}\mathcal{L}_{std}. \quad (8)$$

As shown in Figure 4b, the network with VAF can effectively discover the localized appearance of faces.

3.3. Hierarchy Construction via Transformer-based Parsing Module

The affine transformation and compositions of small subparts make them feasible enough to handle pose and appearance variations. However, they have ambiguous semantics due to their limited expressive power. As shown in Figure 5, the discovered skin template may respond to both forehead and cheekbone and the eyes-related template highlights both left and right eyes. Therefore, if the subparts are assembled with appearance as clues only, there will be ambiguity in the face hierarchy, leading to flaws of the aggregated parts.

We propose a Transformer-based Parsing Module (TPM) to automatically construct the stable subpart-part hierarchy. If we regard each subpart capsule as a visual word, the collection of subpart capsules becomes a sentence, which is very long since we need numbers of subparts (75 in this paper) to handle the complicated appearance of faces. Transformer [30], which is originally proposed to capture the long-range dependencies across words in natural language processing, is naturally extendable to our task. In this paper, we adopt transformer to model the subpart-part relationships, where the subpart capsules $\Theta_k^s : \{p_k^s, \theta_k^s, T_k^s, V_k^s, f_k^s\}$ are sent as input sequences to estimate part capsules $\Theta_m^p : \{T_m^p, V_m^p, r_m^p\}$:

$$\begin{aligned}r_1^p, r_2^p, \dots, r_M^p &= \text{Transformer}(\Theta_1^s, \Theta_2^s, \dots, \Theta_K^s), \\ T_m^p &= \sum_k r_{m_k}^p p_k^s C_k^s \hat{T}_k^s, \\ V_m^p &= \sum_k r_{m_k}^p p_k^s \hat{V}_k^s,\end{aligned}\quad (9)$$

where $r_m^p = [r_{m_1}^p, \dots, r_{m_K}^p]$ describes the probability that subpart capsules belongs to the m -th part. During training, we utilize K-means [23] to generate pseudo subpart-part relationships for the transformer and the \mathcal{L}_{cls} can be




Figure 5. Examples of subparts and their corresponding heatmaps. Some smaller subparts lack semantics due to their ambiguity at different positions.

formulated as:

$$\begin{aligned}\mathcal{L}_{cls} &= -\frac{1}{K} \sum_k \sum_m \tilde{r}_{m_k}^p \log(r_{m_k}^p), \\ \tilde{r}_m^p &= \text{K-means}\left(\left[\hat{V}_k^s + t_{x_k}^s; \hat{V}_k^s + t_{y_k}^s\right]\right),\end{aligned}\quad (10)$$

where $t_{x_k}^s, t_{y_k}^s$ are the translation parameters derived from the pose params θ_k^s . Besides, we also assign a silhouette loss \mathcal{L}_{silh} to constrain the shape consistency across samples:

$$\mathcal{L}_{silh} = \frac{1}{M} \sum_m \|V_m^p - \tilde{V}_m^p\|_2, \quad (11)$$

where \tilde{V}_m^p is the averaged visibility map for the mini-batch.

As for subpart assembling, since each subpart capsule should only belong to one part capsule, we penalize the sparsity of compositional relations with \mathcal{L}_{rela} :

$$\mathcal{L}_{rela} = \frac{1}{K} \sum_k \left[-\sum_m r_{m_k}^p \log(r_{m_k}^p) \right]. \quad (12)$$

The proposed TPM incorporates both appearance and geometric clues for subparts assembling so that the generated parts have more prominent semantics. The overall loss function for hierarchy construction can be formulated as:

$$\mathcal{L}_{part} = \lambda_{cls}\mathcal{L}_{cls} + \lambda_{silh}\mathcal{L}_{silh} + \lambda_{rela}\mathcal{L}_{rela}, \quad (13)$$

where $\lambda_{cls}, \lambda_{silh}$ and λ_{rela} are the hyper-parameters to combine different loss functions.

4. Experiments

4.1. Implementation Details

HP-Capsule contains two sub-networks: the subpart capsule autoencoder with VAF and the Transformer-based

Parsing Module. We replace the four convolutional layers in PCAE [17] with two Residual Blocks [11] as the capsule autoencoder. For the transformer, we adopt the slot attention [21] with learnable slots as encoder and use MLP as decoder, where the iteration of slot attention is set to 3. We extract the foreground mask by SSFNet [41] to concentrate on the part discovery. For fair comparisons, we do the same operation for other unsupervised face segmentation methods that required foreground masks.

During training, the capsule autoencoder is first trained to converge and then refined together with TPM. For optimization, we use Adam optimizer [16] with 10^{-4} learning rate. All experiments are implemented on Pytorch with a single NVIDIA Tesla M40 GPU.

4.2. Datasets

BP4D. We evaluate our model on the BP4D [42] from FG3D dataset [45]. After 3D augmentation by [45], it contains 23,359 images from 41 subjects with various expressions and poses. We randomly choose 90% for training and 10% for testing.

Multi-PIE. The Multi-PIE dataset [9] contains 337 subjects with various expressions. The images are captured under 15 different views spaced in 15° intervals. In this paper, we choose the images from the middle 5 cameras with flash to avoid extreme poses, which contains 8,974 images. In the experiments, Multi-PIE is used to test the generalization ability. We split the dataset as 90% for training and 10% for testing, as the evaluation of unsupervised face segmentation also needs samples to train the fitting model.

4.3. Evaluation Metrics

The Normalized Concentrated Distance (NCD) is used in the ablation study to estimate the concentration of parts, which is formulated as:

$$\text{NCD} = \frac{1}{N} \sum_n \sum_{i,j} \|(i,j) - (c_i^n, c_j^n)\|_2 \cdot S_{n_{i,j}} / z_n, \quad (14)$$

$$(c_i^n, c_j^n) = \left(\sum_{i,j} i \cdot S_{n_{i,j}} / z_n, \sum_{i,j} j \cdot S_{n_{i,j}} / z_n \right),$$

where S_n is the covered region of n -th part, (c_i^n, c_j^n) is the centroid of the n -th part, $z_n = \sum_{i,j} S_{n_{i,j}}$ is for normalization.

The Normalized Mean Error (NME) of landmarks is used as an alternative way to evaluate the parsing quality, which is formulated as:

$$\text{NME} = \frac{1}{N} \sum_n \frac{\|v_n - v_n^*\|_2}{d}, \quad (15)$$

where N is the number of landmarks, v_n is the landmark predicted from the segmentation results, v_n^* is the

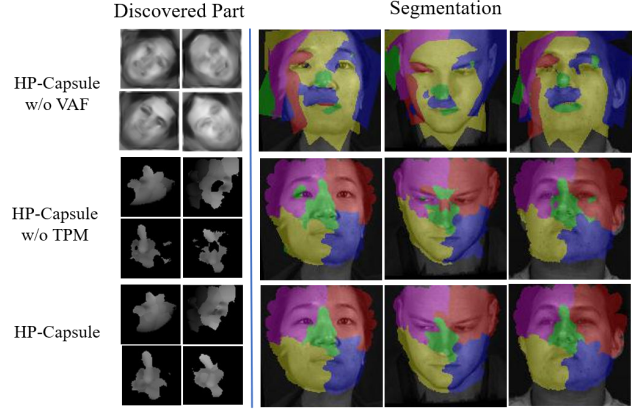


Figure 6. The qualitative ablation study of HP-Capsule. We could see that VAF helps discover effective subparts and TPM improves the semantics of parts.

Table 1. The quantitative ablation study on BP4D. The results show the importance of VAF and TPM for hierarchical face part discovery.

VAF	TPM	NCD	NME _{DL} (%)
		30.72	-
✓		19.27	6.21
✓	✓	18.79	6.10

ground truth and d is the inter-ocular distance. Following SCOPS [15], NME_L uses the centroids of segmentation maps as landmarks and converts them to the human-annotated landmarks by a linear mapping. During the test, we map the predicted 5 landmarks to 5 and 68 ground-truth landmarks separately. However, the centroids of the segmentation maps are too coarse to measure the parsing manners. In this paper, we propose NME_{DL} to evaluate the semantic consistency on a detailed level. NME_{DL} uses a very shallow network to predict the landmarks directly from the segmentation maps, which contains only one Residual Block and one linear layer.

4.4. Ablation Study

We perform the ablation study to show the importance of the VAF for subpart discovery and the TPM for hierarchy construction.

Figure 6 shows the visualization results on face part discovery and segmentation. The first row shows that the original capsule network can not handle the faces, where the discovered parts are holistic and cover almost the whole face. In the second row, with the VAF, the network can capture effective local parts, as the regions with lower visibility are suppressed during training. However, the subparts have ambiguous semantics due to their limited

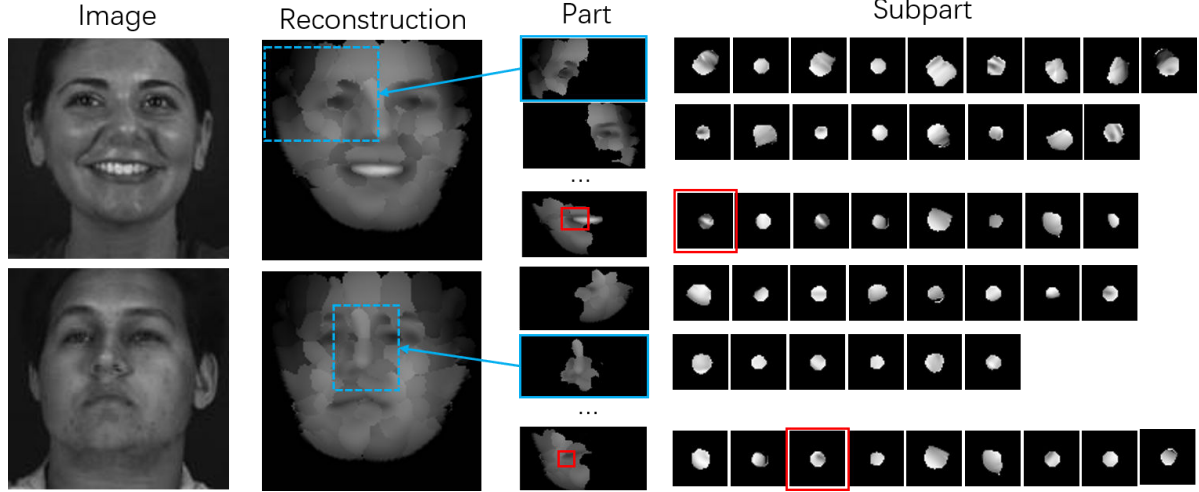


Figure 7. The hierarchical face parts discovered by HP-Capsule. For each input, HP-Capsule automatically selects a set of subparts to describe current object (marked with red boxes) and aggregates them to get parts with more prominent semantics (marked with blue dot boxes), constructing a bottom-up hierarchy.

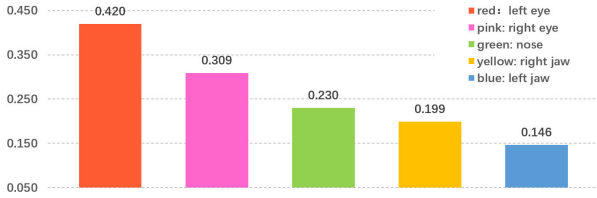


Figure 8. The importance of each region for face recognition. The color of columns is identical to segmentation maps. The regions with eyes are more important for recognition.

expressive power, demonstrated in Figure 5. Therefore, if only using the pseudo subpart-part relationships for hierarchy construction (Eqn. 10), the aggregated parts will suffer serious flaws and become ambiguous in semantics, shown in the second row. After introducing the TPM for refinement, the semantics of parts is improved. The importance of VAF and TPM is further validated by the quantitative evaluation in Table 1.

4.5. Analysis on the Exploited Face Hierarchy

For each input, HP-Capsule automatically activates a subset of subpart capsules and aggregates them to higher-level part capsules, constructing a bottom-up hierarchy. Figure 7 shows the hierarchical face parts discovered by HP-Capsule. It can be seen that HP-Capsule almost reconstructs the original images and parses them into five parts: nose, left side of mouth and cheek, right side of mouth and cheek, the left upper face containing left eye and left ear, and the right upper face. The last column shows the corresponding subparts of each part. Through the templates

defined in capsules, the discovered concepts can be directly visualized without any decoding operations, which makes them easier for us to understand.

We also design a toy experiment to show which part capsule is important for face recognition. We assign each subpart capsule a trainable non-negative scalar w_k as attention weight and send the weighted subpart capsule parameters $w_k \Theta_k^s$ to a linear classifier for face recognition: $y = \text{Linear}(w_1 \Theta_1^s, \dots, w_K \Theta_K^s)$. The classifier is trained with softmax and L1 penalization for the sparsity of w_k . The averaged weights for each part are shown in Figure 8. It can be seen that eye regions are most important for face recognition, which is consistent with the work of Williford *et al.* [32] that shows eyes and nose contain more discriminative features for face recognition.

4.6. Comparison on Unsupervised Face Segmentation

The covered regions of parts discovered by HP-Capsule can be used for the unsupervised face segmentation task, enabling the comparison with other state-of-the-art methods.

Methods. The unsupervised face segmentation from unlabeled images is a challenging task that has not been well explored. DFF [5] proposes to use non-negative matrix factorization upon the CNN features to discover the semantic concepts, which needs to optimize on the whole datasets during inference to keep semantic consistency. SCOPS [15] incorporates the invariance between TPS transformation as clues and proposes a framework with several loss functions for unsupervised segmentation. However, their method mainly relies on the concentration loss, which tends to cut the images into the same region regardless of the face poses.



Figure 9. The qualitative comparison of unsupervised face segmentation on BP4D. HP-Capsule shows better semantic consistency across samples.



Figure 10. The qualitative comparison of unsupervised face segmentation on Multi-PIE. HP-Capsule shows better stability when generalized to another dataset.

Table 2. The quantitative comparison of unsupervised face segmentation on BP4D. $NME_L(\%)$ and $NME_{DL}(\%)$ evaluate the semantic consistency of landmarks.

Method	NME_L		NME_{DL}
	$N_{GT} = 5$	$N_{GT} = 68$	
DFF [5]	18.85	18.62	12.26
SCOPS [15]	9.10	9.67	6.74
HP-Capsule	8.81	9.11	6.10

Evaluation on BP4D. From Figure 9, it can be seen that without any spatial transformation as clues, HP-Capsule still outperforms other methods with better semantic consistency. As shown in the second row, the results of SCOPS sometimes take the right eye as the right green part, sometimes take it as the top pink part or even split it from the middle. In comparison, HP-Capsule keeps the same parsing manners under different poses. This is because the shared templates of HP-Capsule encode the frequently observed patterns and the intrinsic geometric clues are captured for better image reconstruction. Results in Table 2 also validate the effectiveness of our method.

Table 3. The quantitative comparison of unsupervised face segmentation on Multi-PIE. $NME_L(\%)$ and $NME_{DL}(\%)$ evaluate the semantic consistency of landmarks.

Method	NME_L		NME_{DL}
	$N_{GT} = 5$	$N_{GT} = 68$	
DFF [5]	21.88	20.21	17.35
SCOPS [15]	16.15	15.25	13.54
HP-Capsule	12.34	12.30	11.57

Evaluation on Multi-PIE. To further investigate the generalization of methods, we employ Multi-PIE for cross-data evaluation. In Figure 10, we can see that the SCOPS cannot separate the foreground well even it has already been trained with the ground-truth masks of BP4D. Different from previous methods, HP-Capsule learns to estimate face parts with explicit semantics and generalizes well on the Multi-PIE dataset, which is also validated in Table 3.

5. Discussion and Conclusion

In this paper, we propose the HP-Capsule for unsupervised face part discovery. By incorporating the Visibility Activation Function and the Transformer-based Parsing Module, the network successfully discovers face parts and constructs the face hierarchy from the unlabeled images. This work extends the application of capsule networks from digits to human faces and provides an insight into the visual perception mechanism of networks. Besides, HP-Capsule gives unsupervised face segmentation results by the covered regions of part capsules, enabling the comparison with other state-of-the-art methods. Experiments on BP4D and Multi-PIE show the superior performance of our method.

Limitation. HP-Capsule presents the face hierarchy with a set of visualizable templates, whose representation power is limited by the number of the templates. Improving the representation ability of templates while maintaining their interpretability remains a challenging task and deserves further studies in future work.

Broader Impacts. The proposed method is a generative model that discovers face hierarchy based on the learned statistics of the training dataset, which will reflect biases in those data, including ones with negative societal impacts. These issues are worth further consideration if the capsules are utilized to generate images.

Acknowledgement. This work was supported in part by the National Key Research & Development Program (No. 2020YFC2003901), Chinese National Natural Science Foundation Projects #62176256, #61876178, #61976229, #62106264, the Youth Innovation Promotion Association CAS (#Y2021131) and the InnoHK program.

References

- [1] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6541–6549, 2017. 3
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3
- [3] Tianshui Chen, Liang Lin, Riquan Chen, Yang Wu, and Xiaonan Luo. Knowledge-embedded representation learning for fine-grained image recognition. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 627–634, 2018. 2
- [4] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16794–16804, 2021. 3
- [5] Edo Collins, Radhakrishna Achanta, and Sabine Susstrunk. Deep feature factorization for concept discovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–352, 2018. 3, 7, 8
- [6] Tao Feng, Stan Z Li, Heung-Yeung Shum, and HongJiang Zhang. Local non-negative matrix factorization as a visual representation. In *Proceedings 2nd International Conference on Development and Learning. ICDL 2002*, pages 178–183. IEEE, 2002. 3
- [7] Qinzhe Gao, Bin Wang, Libin Liu, and Baoquan Chen. Unsupervised co-part segmentation through assembly. In *International Conference on Machine Learning*, 2021. 3
- [8] Abel Gonzalez-Garcia, Davide Modolo, and Vittorio Ferrari. Do semantic parts emerge in convolutional neural networks? *International Journal of Computer Vision*, 126(5):476–494, 2018. 3, 4
- [9] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and vision computing*, 28(5):807–813, 2010. 6
- [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [12] Geoffrey Hinton. Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science*, 3(3):231–250, 1979. 1
- [13] Geoffrey E Hinton, Alex Krizhevsky, and Sida D Wang. Transforming auto-encoders. In *International conference on artificial neural networks*, pages 44–51. Springer, 2011. 2
- [14] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with em routing. In *International conference on learning representations*, 2018. 2
- [15] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 869–878, 2019. 3, 4, 5, 6, 7, 8
- [16] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, page arXiv:1412.6980, Dec. 2014. 6
- [17] AR Kosiorek, S Sabour, YW Teh, GE Hinton, and Miss Jo STAFFORD-TOLLEY. Stacked capsule autoencoders. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 4, 6
- [18] Jinpeng Lin, Hao Yang, Dong Chen, Ming Zeng, Fang Wen, and Lu Yuan. Face parsing with roi tanh-warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5654–5663, 2019. 2
- [19] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. Unsupervised part segmentation through disentangling appearance and shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8355–8364, 2021. 3
- [20] Yinglu Liu, Hailin Shi, Hao Shen, Yue Si, Xiaobo Wang, and Tao Mei. A new dataset and boundary-attention semantic segmentation for face parsing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11637–11644, 2020. 2
- [21] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *NeurIPS*, 2020. 6
- [22] David Marr and Herbert Keith Nishihara. Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 200(1140):269–294, 1978. 1
- [23] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011. 5
- [24] Ruijie Quan, Xuanyi Dong, Yu Wu, Linchao Zhu, and Yi Yang. Auto-reid: Searching for a part-aware convnet for person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3750–3759, 2019. 2
- [25] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3859–3869, 2017. 2
- [26] Sara Sabour, Andrea Tagliasacchi, Soroosh Yazdani, Geoffrey Hinton, and David J Fleet. Unsupervised part representation by flow capsules. In *International Conference on Machine Learning*, pages 9213–9223. PMLR, 2021. 2, 3
- [27] Manish Singh and Donald D Hoffman. Part-based representations of visual shape and implications for visual

- cognition. In *Advances in psychology*, volume 130, pages 401–459. Elsevier, 2001. 1
- [28] Maneet Singh, Shruti Nagpal, Richa Singh, and Mayank Vatsa. Dual directed capsule network for very low resolution image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 340–349, 2019. 2
- [29] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10052–10062, 2021. 3
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 5
- [31] Zhenzhen Weng, Mehmet Giray Ogut, Shai Limonchik, and Serena Yeung. Unsupervised discovery of the long-tail in instance segmentation using hierarchical self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2603–2612, 2021. 3
- [32] Jonathan R Williford, Brandon B May, and Jeffrey Byrne. Explainable face recognition. In *European Conference on Computer Vision*, pages 248–263. Springer, 2020. 7
- [33] Jinlin Wu, Yang Yang, Zhen Lei, Jinqiao Wang, Stan Z Li, Prayag Tiwari, and Hari Mohan Pandey. An end-to-end exemplar association for unsupervised person re-identification. *Neural Networks*, 129:43–54, 2020. 3
- [34] Jinlin Wu, Yang Yang, Hao Liu, Shengcai Liao, Zhen Lei, and Stan Z Li. Unsupervised graph association for person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8321–8330, 2019. 3
- [35] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2020. 3
- [36] Xixin Wu, Songxiang Liu, Yuwen Cao, Xu Li, Jianwei Yu, Dongyang Dai, Xi Ma, Shoukang Hu, Zhiyong Wu, Xunying Liu, et al. Speech emotion recognition using capsule networks. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6695–6699. IEEE, 2019. 2
- [37] Yuan Xie, Tianshui Chen, Tao Pu, Hefeng Wu, and Liang Lin. Adversarial graph representation adaptation for cross-domain facial expression recognition. In *Proceedings of the 28th ACM international conference on Multimedia*, pages 1255–1264, 2020. 2
- [38] Jingwei Xin, Nannan Wang, Xinrui Jiang, Jie Li, Xinbo Gao, and Zhifeng Li. Facial attribute capsules for noise face super resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12476–12483, 2020. 2
- [39] Zhenjia Xu, Zhijian Liu, Chen Sun, Kevin Murphy, William T Freeman, Joshua B Tenenbaum, and Jiajun Wu. Modeling parts, structure, and system dynamics via predictive learning. In *International Conference on Learning Representations*, 2019. 3
- [40] Xiaomei Zhang, Yingying Chen, Bingke Zhu, Jinqiao Wang, and Ming Tang. Part-aware context network for human parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8971–8980, 2020. 2
- [41] Xiaomei Zhang, Yingying Chen, Bingke Zhu, Jinqiao Wang, and Ming Tang. Semantic-spatial fusion network for human parsing. *Neurocomputing*, 402:375–383, 2020. 6
- [42] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, Peng Liu, and Jeffrey M Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014. 6
- [43] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Pscanet: Point-wise spatial attention network for scene parsing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 267–283, 2018. 2
- [44] Yongheng Zhao, Tolga Birdal, Haowen Deng, and Federico Tombari. 3d point capsule networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1009–1018, 2019. 2
- [45] Xiangyu Zhu, Fan Yang, Di Huang, Chang Yu, Hao Wang, Jianzhu Guo, Zhen Lei, and Stan Z Li. Beyond 3dmm space: Towards fine-grained 3d face reconstruction. In *European Conference on Computer Vision*, pages 343–358. Springer, 2020. 6