

1DFormer: A Transformer Architecture Learning 1D Landmark Representations for Facial Landmark Tracking

Shi Yin^{1,2}, Shijie Huang¹, Shangfei Wang³, Jinshui Hu^{1*}, Tao Guo¹, Bing Yin¹, Baocai Yin¹ and Cong Liu^{1*}

¹iFLYTEK Research

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³University of Science and Technology of China

shiyin@iai.ustc.edu.cn, sjhuang6@iflytek.com, sfwang@ustc.edu.cn, jshu@iflytek.com, taoguo@iflytek.com, bingyin@iflytek.com, bcyin@iflytek.com, congliu2@iflytek.com

Abstract

Recently, heatmap regression methods based on 1D landmark representations have shown prominent performance on locating facial landmarks. However, previous methods ignored to make deep explorations on the good potentials of 1D landmark representations for sequential and structural modeling of multiple landmarks to track facial landmarks. To address this limitation, we propose a Transformer architecture, namely 1DFormer, which learns informative 1D landmark representations by capturing the dynamic and the geometric patterns of landmarks via token communications in both temporal and spatial dimensions for facial landmark tracking. For temporal modeling, we propose a confidence-enhanced multi-head attention mechanism with a recurrently token mixing strategy to adaptively and robustly embed long-term landmark dynamics into their 1D representations; for structure modeling, we design intra-group and inter-group geometric encoding mechanisms to encode the component-level as well as global-level facial structure patterns as a refinement for the 1D representations of landmarks through token communications in the spatial dimension via 1D convolutional layers. Experimental results on the 300VW and the TF databases show that 1DFormer successfully models the long-range sequential patterns as well as the inherent facial structures to learn informative 1D representations of landmark sequences, and achieves state-of-the-art performance on facial landmark tracking. Codes of our model are available in the supplementary materials.

1 Introduction

Tracking facial landmarks from a video stream [Wu and Ji, 2019; Chrysos *et al.*, 2018] is a fundamental task for human-centered applications, such as human-computer interaction, emotion analysis, and face recognition [Wang *et al.*, 2021b], while remaining unsolved in the challenging “in-the-wild” scenarios. Recently, 1D heatmap regression methods, which are built based on 1D landmark representations, i.e., the light-weight 1D feature vectors and heatmaps representing the marginal distribution of every landmark on each axis as depicted in Figure 1, have achieved prominent performance for landmark localization of human faces [Yin *et al.*, 2020; Xiong *et al.*, 2020; Kang *et al.*, 2023], human bodies [Moon and Lee, 2020; Li *et al.*, 2022c; Chi *et al.*, 2022] and objects [Liu *et al.*, 2022]. Compared to the 2D landmark representations [Newell *et al.*, 2016; Bulat and Tzimiropoulos, 2017; Wang *et al.*, 2021a] with quadratic spatial complexity, the 1D representations, with a linear spatial complexity, could achieve higher resolution on each axis and enough feature channels under the limited machine memory. Thus, these methods based on 1D representations could capture more details of the spatial patterns of landmarks and bring a lower quantization error as well as a better accuracy on locating facial landmarks under the same hardware condition. However, previous methods have not fully explored the good potentials of 1D landmark representations for temporal sequence modeling, which is critical for the facial landmark tracking task. What’s more, all of these methods predicted the 1D representation of each landmark in a loosely coupled manner, lacking considerations of structural modeling upon multiple landmarks. These limitations caused the accuracy bottleneck of these methods.

To address these, as shown in Figure 2, we propose a Transformer-based 1D representation learning method, i.e., 1DFormer, to effectively capture the temporal and structural patterns of multiple landmarks for facial landmark tracking. The clusters of 1D landmark representations, including the hidden features as well as the output heatmaps, are in the form of a sequence of 1D light-weight vectors, which are the ideal input and output patterns of Transformer. Thus, we build our method based on the paradigm of Transformer and develop a Transformer architecture with both temporal modeling and

*Corresponding Authors

The links of supplementary materials: https://drive.google.com/file/d/1J0q5eY4R_5iQ-3Wa8ZaHQgFJtjeZwK-X/view?usp=drive_link, also available at: <https://pan.baidu.com/s/1MPpD4S7h2x5rcUhQo5nF2g>, extraction code: ml75

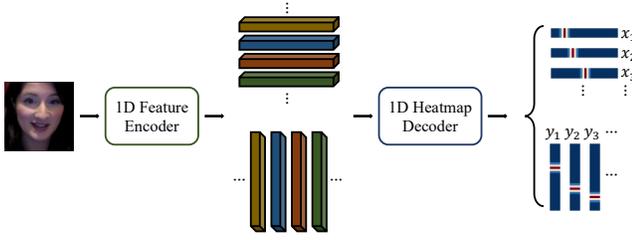


Figure 1: An illustration of 1D heatmap regression methods, which are built upon 1D landmark representations, including 1D feature vectors and heatmaps.

structural modeling techniques. For temporal modeling, we propose a confidence-enhanced multi-head attention mechanism with a recurrently token mixing strategy to integrate long-range temporal information from past frames to enhance the features of the current time step through efficient window sliding and information delivery. The attention mechanism combines the temporal correlations learned from the key-query mechanism as well as the confidence score of feature qualities as the integrated attention weights, robustly guiding the fusion of features from different time steps despite challenging imaging conditions. For structural modeling, to make full use of the intrinsic geometry patterns of faces for landmark tracking, we design an intra-group and an inter-group geometry encoding mechanisms to embed component-level as well as global-level facial structure patterns as a refinement for the 1D representations of the landmark sequence. It is expected that these mechanisms can effectively extract the structural patterns of multiple landmarks, and meanwhile, keep the clear semantic of each landmark. For that purpose, we adopt 1D convolutional layers instead of the vanilla attention layers for token communications in the spatial dimension. Through the structural modeling mechanisms, the learned facial structures are helpful to correct the texture disturbances caused by occlusions or uneven illumination conditions, and further improve the landmark tracking performance for “in-the-wild” scenarios.

The main contributions of this paper can be summarized as the following. First, as far as we know, we are the first to fit the prominent Transformer paradigm on learning informative 1D representations of facial landmark sequences for facial landmark tracking. Second, we propose a new Transformer architecture, namely 1DFormer, with both temporal modeling and structural modeling mechanisms, to explore the good potentials of 1D representation learning on modeling the long-term sequential as well as the geometric patterns of facial landmarks. Third, as demonstrated by experiments on the 300VW database and the TF database, the tracker based on 1DFormer achieves state-of-the-art performance for facial landmark tracking with a significant improvement of accuracy and stability performance compared to the related works.

2 Related Work

As the mainstream of key point detection and tracking, heatmap regression methods explicitly modeled the spatial distributions of landmarks by learning to predict the proba-

bility intensity that a landmark appears in each position of the 2D plane or 3D space. Through explicitly capturing the landmark distributions, these methods achieved a better spatial generalization performance on facial landmark detection and tracking compared to the conventional coordinate regression methods which directly regressed the values of landmark coordinates [Xiong and De la Torre, 2013; Liu *et al.*, 2018; Simonyan and Zisserman, 2014; Yin *et al.*, 2019; Li *et al.*, 2013; Wu *et al.*, 2014; Zhu *et al.*, 2022a]. Heatmap regression methods can be further divided into two categories: 2D heatmap regression and 1D heatmap regression.

2D heatmap regression methods [Newell *et al.*, 2016; Peng *et al.*, 2016; Bulat and Tzimiropoulos, 2017; Wang *et al.*, 2021a; Li *et al.*, 2022a; Li *et al.*, 2022b] modeled 2D joint distributions, typically assumed as 2D Gaussian distributions, on the x and y axes in both hidden and output stages for each landmark. These approaches excelled in spatial generalization but faced challenges due to the high spatial complexity, i.e., $\mathcal{O}(L^2)$ where L is the resolution of 2D features and heatmaps, requiring down-sampling of feature and heatmap resolutions under limited machine space. This led to precision loss and quantization errors, thus limiting the effectiveness in landmark detection tracking. To address this, 1D heatmap regression methods [Yin *et al.*, 2020; Xiong *et al.*, 2020; Moon and Lee, 2020; Li *et al.*, 2022c; Chi *et al.*, 2022; Liu *et al.*, 2022] replaced some 2D landmark representations of high spatial complexity with 1D features and heatmaps. These methods, assuming 1D Gaussian distribution as the marginal distribution of landmarks on each axis, used strided CNN [Yin *et al.*, 2020], MLP [Li *et al.*, 2022c], or pooling [Moon and Lee, 2020] to compress image representations along one axis while enhancing resolution along the other. With a feature resolution of L , the spatial complexity is reduced to $\mathcal{O}(L)$, allowing higher resolutions per axis and more feature channels to capture fine-grained details of landmark distributions. This brings improved landmark localization performance compared to both coordinate and 2D heatmap regression methods.

However, previous 1D heatmap regression methods, mostly detection-based [Xiong *et al.*, 2020; Moon and Lee, 2020; Li *et al.*, 2022c; Chi *et al.*, 2022; Liu *et al.*, 2022; Kang *et al.*, 2023], lacked depth in time series modeling crucial for landmark tracking. To the best of our knowledge, only Yin *et al.* [2020] applied 1D representations to model temporal sequences of landmarks. However, their sequential encoder was quite simple with fixed preset weights for fusing features from the past frames, without a dynamic weight assignment mechanism to adaptively model the temporal correlations and confidences of feature qualities from different frames. These preset weights also decreased quickly with the exponential decay of time, causing a sub-optimal performance on capturing long-term temporal dependencies from historical frames. What’s more, these methods ignored to develop an effective interaction mechanism to capture the structural patterns of multiple landmarks. To address these issues, we propose a new Transformer architecture adept at long-range temporal modeling and geometric modeling to enhance 1D representation learning for facial landmark tracking.

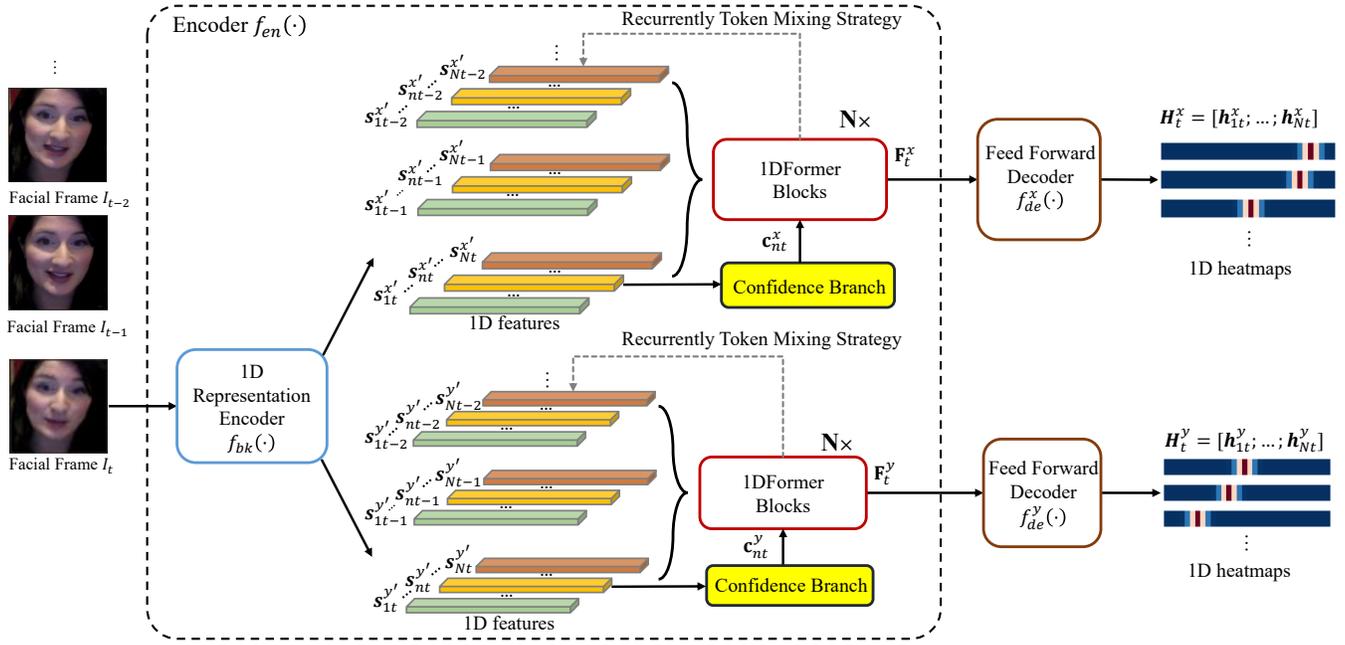


Figure 2: An architectural overview of the proposed facial landmark tracking method, i.e., 1DFormer.

3 Method

The proposed facial landmark tracker consists of an 1D feature encoder $f_{en}(\cdot)$ as well as two 1D heatmap decoders $f_{de}^x(\cdot)$ and $f_{de}^y(\cdot)$. The video to be extracted facial landmarks is denoted as $\mathbf{I}_{1:T} = [\mathbf{I}_1; \mathbf{I}_2, \dots; \mathbf{I}_T]$, where $\mathbf{I}_t \in \mathbb{R}^{H \times W \times 3}$ ($1 \leq t \leq T$) is the image of the t th frame. $f_{en}(\cdot)$ encodes $\mathbf{I}_{1:T}$ as a cluster of features representing the marginal distributions of landmarks on the x and y axes for each time step:

$$\mathbf{F}_1^x, \mathbf{F}_1^y, \mathbf{F}_2^x, \mathbf{F}_2^y, \dots, \mathbf{F}_T^x, \mathbf{F}_T^y = f_{en}(\mathbf{I}_{1:T}; \theta_{en}) \quad (1)$$

where θ_{en} is the parameters of $f_{en}(\cdot)$, $\mathbf{F}_t^x = [\mathbf{f}_{1t}^x; \mathbf{f}_{2t}^x; \dots; \mathbf{f}_{Nt}^x] \in \mathbb{R}^{N \times L}$ and $\mathbf{F}_t^y = [\mathbf{f}_{1t}^y; \mathbf{f}_{2t}^y; \dots; \mathbf{f}_{Nt}^y] \in \mathbb{R}^{N \times L}$ respectively denote the 1D representations on x axis and the y axis at the t th time step. Here N is the number of the pre-defined landmarks, \mathbf{f}_{nt}^x ($1 \leq n \leq N$) $\in \mathbb{R}^{1 \times L}$ and $\mathbf{f}_{nt}^y \in \mathbb{R}^{1 \times L}$ denote the 1D features for the n th landmark at the t th time step. Based on \mathbf{F}_t^x and \mathbf{F}_t^y , the fully connected feed forward 1D heatmap decoders $f_{de}^x(\cdot)$ and $f_{de}^y(\cdot)$, respectively predict two groups of 1D heatmaps, i.e., $\mathbf{h}_{nt}^x = f_{de}^x(\mathbf{f}_{nt}^x) \in \mathbb{R}^{1 \times D}$ and $\mathbf{h}_{nt}^y = f_{de}^y(\mathbf{f}_{nt}^y) \in \mathbb{R}^{1 \times D}$, for the two axes, correspondingly. The x and y coordinates of the landmark can be obtained from the peak positions of \mathbf{h}_{nt}^x and \mathbf{h}_{nt}^y , respectively. In this paper, we propose a new design of $f_{en}(\cdot)$, which is composed of a backbone encoder extracting 1D representations from each individual frame image, as well as a Transformer, which we called as 1DFormer, to extract long-range sequential patterns and geometric structures of facial landmarks to refine their 1D representations.

3.1 Backbone Encoder

Before processed by 1DFormer, we first extract 1D representations for every landmark from each individual frame image

through a backbone 1D representation encoder, which is composed of a FAN network [Bulat and Tzimiropoulos, 2017] cascaded with two groups of strided CNNs [Yin *et al.*, 2020] to encode 1D representations, i.e., \mathbf{s}_{nt}^x ($1 \leq n \leq N, 1 \leq t \leq T$) $\in \mathbb{R}^{1 \times L}$ and \mathbf{s}_{nt}^y , on the two axes, respectively.

3.2 Temporal Modeling Mechanism of 1DFormer

We take attention mechanism to embed the dynamic patterns of the landmark sequence into \mathbf{s}_{nt}^x ($1 \leq n \leq N, 1 \leq t \leq T$) and \mathbf{s}_{nt}^y , respectively, by integrating the 1D landmark representations from different time steps. The inner structures of our attention modules will be discussed later. Here we firstly present a recurrently token mixing strategy arranging the inputs and outputs of the attention modules to capture the long-range sequential patterns of facial landmarks from the video stream, as shown in Eq. (2):

$$\begin{aligned} \mathbf{s}_{nt}^{x'} &= \text{Attention}(\mathbf{s}_{nt}^x, \mathbf{s}_{n(t-1)}^{x'}, \dots, \mathbf{s}_{n(t-W+1)}^{x'}) \\ \mathbf{s}_{nt}^{y'} &= \text{Attention}(\mathbf{s}_{nt}^y, \mathbf{s}_{n(t-1)}^{y'}, \dots, \mathbf{s}_{n(t-W+1)}^{y'}) \end{aligned} \quad (2)$$

where $\mathbf{s}_{nt}^{x'}$ and $\mathbf{s}_{nt}^{y'}$ ($1 \leq n \leq N, 1 \leq t \leq T$) are the features integrated with temporal information of the video through mixing temporal tokens, i.e., $\mathbf{s}_{nt}^x, \mathbf{s}_{n(t-1)}^{x'}, \dots, \mathbf{s}_{n(t-W+1)}^{x'}$ or $\mathbf{s}_{nt}^y, \mathbf{s}_{n(t-1)}^{y'}, \dots, \mathbf{s}_{n(t-W+1)}^{y'}$. As a computationally acceptable way, instead of taking the features from all of the past frames as the input, we just send the features from a temporal neighborhood window of W frames into the attention modules to refine the 1D features of landmarks with sequential patterns. Since the memories of temporal patterns are delivered with the sliding of the window from the past windows to the current window, $\mathbf{s}_{nt}^{x'}$ and $\mathbf{s}_{nt}^{y'}$ can still keep a long memory of the landmark dynamics.

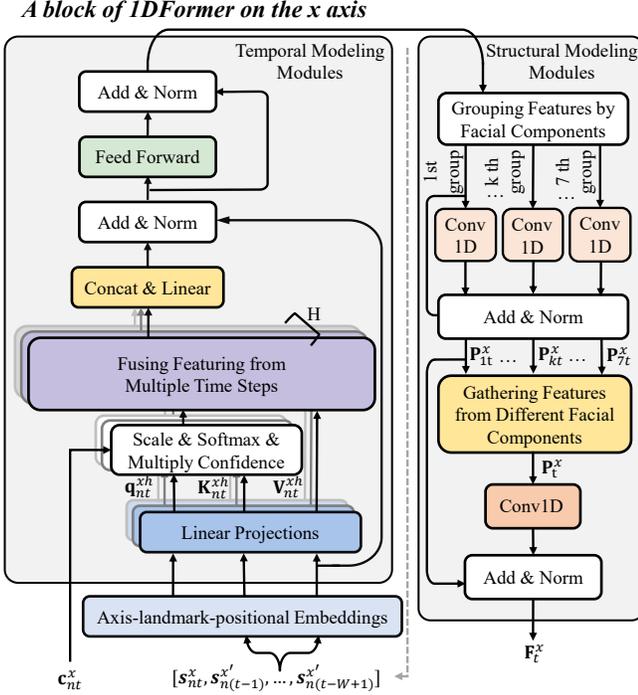


Figure 3: The internal architecture of a basic block of 1DFormer on the x axis. The architecture of basic block on the y axis is the same as the x axis.

We extend the conventional positional embedding of Transformer as axis-landmark-positional (alp) embeddings which distinguish not only different position indexes but also different axes and landmarks, to make the attention modules aware of the discrepancy of marginal distributions on different axes, the semantic distinctions among different facial landmarks, as well as the relative position of each feature vector from a window. These embeddings are denoted as \mathbf{e}_{nw}^x and \mathbf{e}_{nw}^y ($1 \leq n \leq N, 1 \leq w \leq W$), respectively, where the superscript (x or y) denotes the axis the 1D feature belongs to, subscript n is the landmark index, w is the relative position index of an input feature to the whole input window with W frames. The input feature sequences for attention modules are added with the axis-landmark-positional embeddings in the way like: $\mathbf{s}_{nt}^x := \mathbf{s}_{nt}^x + \mathbf{e}_{n1}^x, \mathbf{s}_{n(t-w+1)}^x := \mathbf{s}_{n(t-w+1)}^x + \mathbf{e}_{nw}^x$ ($1 < w \leq W$), $\mathbf{s}_{nt}^y := \mathbf{s}_{nt}^y + \mathbf{e}_{n1}^y$, $\mathbf{s}_{n(t-w+1)}^y := \mathbf{s}_{n(t-w+1)}^y + \mathbf{e}_{nw}^y$ ($1 < w \leq W$). Here \mathbf{e}_{nw}^x and \mathbf{e}_{nw}^y are randomly initialized, then optimized with the tracker.

Next, we present the inner structure of our attention mechanism, which specifically is a confidence-enhanced multi-head attention mechanism, to dynamically determine the attention weights of features from different time steps according to their temporal correlations as well as their qualities influenced by the imaging conditions. Typically, the attention weights can be calculated by the scaled dot-product opera-

tions between the query and the key matrices:

$$\begin{aligned} \mathbf{q}_{nt}^{xh} &= \mathbf{s}_{nt}^x \mathbf{W}_q^{xhT}, & \mathbf{q}_{nt}^{yh} &= \mathbf{s}_{nt}^y \mathbf{W}_q^{yhT} \\ \mathbf{K}_{nt}^{xh} &= \mathbf{I}_{nt}^x \mathbf{W}_k^{xhT}, & \mathbf{K}_{nt}^{yh} &= \mathbf{I}_{nt}^y \mathbf{W}_k^{yhT} \\ \mathbf{a}_{nt}^{xh} &= \text{sf}\left(\frac{\mathbf{q}_{nt}^{xh} \mathbf{K}_{nt}^{xhT}}{\sqrt{d_h}}\right), & \mathbf{a}_{nt}^{yh} &= \text{sf}\left(\frac{\mathbf{q}_{nt}^{yh} \mathbf{K}_{nt}^{yhT}}{\sqrt{d_h}}\right) \end{aligned} \quad (3)$$

where $\text{sf}(\cdot)$ denotes the *Softmax* function; $\mathbf{I}_{nt}^x = [\mathbf{s}_{nt}^x; \mathbf{s}_{n(t-1)}^x; \dots; \mathbf{s}_{n(t-W+1)}^x] \in \mathbb{R}^{W \times L}$ and \mathbf{I}_{nt}^y denote the feature sequence for the t th input window chunk of the attention module; h denotes the head index of the multi-head attention mechanism; \mathbf{q}_{nt}^{xh} and \mathbf{K}_{nt}^{xh} are the h ($1 \leq h \leq H$) th query vector pair to get the attention weight vectors, i.e., $\mathbf{a}_{nt}^{xh} \in \mathbb{R}^{1 \times W}$ and $\mathbf{a}_{nt}^{yh} \in \mathbb{R}^{1 \times W}$, for the t the time step; \mathbf{K}_{nt}^{xh} and $\mathbf{K}_{nt}^{yh} \in \mathbb{R}^{W \times d_h}$ are the h th key matrix pair; \mathbf{W}_q^{xh} , \mathbf{W}_q^{yh} , \mathbf{W}_k^{xh} , and $\mathbf{W}_k^{yh} \in \mathbb{R}^{d_h \times L}$ are the projection matrices. By optimizing the attention mechanism, \mathbf{a}_{nt}^{xh} and \mathbf{a}_{nt}^{yh} can implicitly capture the correlations among features of the temporal sequence: the intensity of the w ($1 \leq w \leq W$) th element of \mathbf{a}_{nt}^{xh} reflects the correlation between \mathbf{s}_{nt}^x and $\mathbf{s}_{n(t-w+1)}^x$.

We extend the conventional attention mechanism with two confidence prediction branches $f_c^x(\cdot)$ and $f_c^y(\cdot)$ to predict the confidence score of feature qualities at each time step. We take \mathbf{s}_{nt}^x and \mathbf{s}_{nt}^y instead of \mathbf{s}_{nt}^x and \mathbf{s}_{nt}^y as the input for the confidence predictors to reflect the imaging quality of each individual video frame. It is expected that, if the texture features around the n th landmark at the t th frame are of good quality with a clear positional discrimination for the landmark, $c_{nt}^x = f_c^x(\mathbf{s}_{nt}^x)$ and $c_{nt}^y = f_c^y(\mathbf{s}_{nt}^y)$ are assigned as a high value and the effects of \mathbf{s}_{nt}^x and \mathbf{s}_{nt}^y on temporal token communications will be encouraged; otherwise, if the image textures of the n th landmark are polluted at the current time step by occlusions, c_{nt}^x and c_{nt}^y are lowered, and \mathbf{s}_{nt}^x as well as \mathbf{s}_{nt}^y will be mitigated to alleviate the effect of feature disturbances. We integrate the weights calculated by the key-query operations with the confidence scores, namely $\mathbf{c}_{nt}^x = [c_{nt}^x, c_{n(t-1)}^x, \dots, c_{n(t-W+1)}^x]$ and $\mathbf{c}_{nt}^y \in \mathbb{R}^{1 \times W}$, as the final attention weights, to model temporal correlations in landmark sequences while maintaining robustness against texture disturbances, as shown in Eq. (4):

$$\mathbf{a}_{nt}^{xh} = \text{sf}\left(\frac{\mathbf{q}_{nt}^{xh} \mathbf{K}_{nt}^{xhT}}{\sqrt{d_h}} \odot \mathbf{c}_{nt}^{xT}\right), \quad \mathbf{a}_{nt}^{yh} = \text{sf}\left(\frac{\mathbf{q}_{nt}^{yh} \mathbf{K}_{nt}^{yhT}}{\sqrt{d_h}} \odot \mathbf{c}_{nt}^{yT}\right) \quad (4)$$

Based on \mathbf{a}_{nt}^{xh} and \mathbf{a}_{nt}^{yh} , the features from different time steps are fused as the temporally refined 1D landmark representations, i.e., $\mathbf{s}_{nt}^{x'}$ and $\mathbf{s}_{nt}^{y'}$, for the n landmark at the t th time step. Since the feature fusion steps are vanilla in Transformer, we just attach its details in the supplementary material due to the page limits.

3.3 Structural Modeling Mechanism of 1DFormer

We embed facial structural patterns to enhance the 1D representation of each landmark, as shown in Figure 3. According to the inherent structure of a human face, its landmarks can be divided into seven groups, i.e., left eyebrow,

right eyebrow, left eye, right eye, nose, mouth, and contour. To encode component-level structural patterns as a refinement for 1D landmark representations, we propose an intra-group geometric encoding mechanism. Specifically, at the $t(1 \leq t \leq T)$ th time step, we stack the features outputted by the attention modules for landmarks of the $k(1 \leq k \leq 7)$ th group as $\mathbf{S}_{kt}^{x'} = \text{stack}(\mathbf{s}_{k(1)t}^{x'}, \mathbf{s}_{k(2)t}^{x'}, \dots, \mathbf{s}_{k(N_k)t}^{x'})$ and $\mathbf{S}_{kt}^{y'} = \text{stack}(\mathbf{s}_{k(1)t}^{y'}, \mathbf{s}_{k(2)t}^{y'}, \dots, \mathbf{s}_{k(N_k)t}^{y'})$, respectively, where the stacking operation gathering multiple 1D representation vectors as multi-channels for convolution, N_k is the total number of landmarks of the k th group, $\mathbf{s}_{k(j)t}^{x'}(1 \leq j \leq N_k)$ and $\mathbf{s}_{k(j)t}^{y'}$ denote the 1D representations produced by Eq. (2) for the j th landmark of the k th group. Here we replace the subscript $n(1 \leq n \leq N)$ of $\mathbf{s}_{nt}^{x'}$ and $\mathbf{s}_{nt}^{y'}$ in Eq. (2) with the subscript $k(j)$ just to mark the switch between the universal index of all landmarks to the relative index within a landmark group. We integrate the component-level structural information into the features of the $k(1 \leq k \leq 7)$ th landmark group via two residual 1D convolutional layers, as shown in Eq. (5):

$$\begin{aligned} \mathbf{P}_{kt}^x &= \text{LN}(\text{Conv1D}(\mathbf{S}_{kt}^{x'}) + \mathbf{S}_{kt}^{x'}) \\ \mathbf{P}_{kt}^y &= \text{LN}(\text{Conv1D}(\mathbf{S}_{kt}^{y'}) + \mathbf{S}_{kt}^{y'}) \end{aligned} \quad (5)$$

where LN denotes the layer normalization operation.

Furthermore, to capture global facial shapes, we adopt an inter-group global geometric encoding mechanism based on the output features from the intra-group geometric encoding module. Let $\mathbf{P}_t^x = \text{stack}(\mathbf{P}_{1t}^x, \mathbf{P}_{2t}^x, \dots, \mathbf{P}_{7t}^x)$, and $\mathbf{P}_t^y = \text{stack}(\mathbf{P}_{1t}^y, \mathbf{P}_{2t}^y, \dots, \mathbf{P}_{7t}^y)$, where the stacking operation gathers the inputs along the channel dimension for convolution. We introduce two residual 1D CNNs to embed global facial structures as a refinement of the 1D representations for each landmark, as shown in Eq. (6):

$$\begin{aligned} \mathbf{F}_t^x &= \text{LN}(\text{Conv1D}(\mathbf{P}_t^x) + \mathbf{P}_t^x) \\ \mathbf{F}_t^y &= \text{LN}(\text{Conv1D}(\mathbf{P}_t^y) + \mathbf{P}_t^y) \end{aligned} \quad (6)$$

where the n th channels of \mathbf{F}_{kt}^x and \mathbf{F}_{kt}^y , denoted \mathbf{f}_{nt}^x and \mathbf{f}_{nt}^y respectively, are the features for the n th landmark refined by global structural patterns. From the proposed geometric encoding modules, facial geometric patterns are captured through feature communications among different facial landmarks, and the learned facial structures help to correct the disturbances on facial appearances caused by occlusions or uneven illumination conditions, and further improve the landmark tracking performance for “in-the-wild” scenarios. It is worth noting that, we take 1D convolutional layers instead of attention layers for token communications in the spatial dimension. The reason lies in that, unlike the temporal tokens which are the features of one certain facial landmark from different time steps, the spatial tokens are features from different landmarks which constitute a structure, and we find that the 1D convolutional layers show a better capability on modeling the general structural patterns among different landmarks while keeping their clear semantic distinguishments. See supplementary materials for a empirical study.

3.4 Optimization

Our method is trained with the following loss functions:

$$\begin{aligned} L_o &= \lambda_h L_h + \lambda_c L_c, \\ L_h &= \sum_{t=1}^T \sum_{n=1}^N (\|\mathbf{h}_{nt}^x - \mathbf{h}_{nt}^{x*}\|_2^2 + \|\mathbf{h}_{nt}^y - \mathbf{h}_{nt}^{y*}\|_2^2), \\ L_c &= \sum_{t=1}^T \sum_{n=1}^N (\|c_{nt}^x - c_{nt}^{x*}\|_2^2 + \|c_{nt}^y - c_{nt}^{y*}\|_2^2) \end{aligned} \quad (7)$$

where λ_h and λ_c are the hyper-parameters determining the weights of each loss in the overall loss function; L_h is the training loss for 1D heatmap regression; \mathbf{h}_{nt}^{x*} and \mathbf{h}_{nt}^{y*} are the heatmap labels, which are 1D Gaussian distributions around the ground truth landmark positions; L_c is the training loss for confidence regression; c_{nt}^{x*} and c_{nt}^{y*} are the confidence labels. Due to the difficulty of manual annotations for confidences on feature qualities, current benchmark databases for facial landmark tracking do not provide such labels. As an alternative, we have to infer confidence labels through a flow-based model [Li *et al.*, 2021] which models landmarks’ positional certainty reflecting the texture qualities, and take its inference results as the pseudo labels of confidences. For a better generalization, we just take these pseudo labels as early guidance of our confidence branches, then fine-tune the confidence branches through the supervised signals from heatmap regression in an end-to-end manner. Formally, denote the total training epochs as E , when $1 \leq e \leq E/2$ we train the tracker with L_o ; when $E/2 \leq e \leq E$ we train the tracker with L_h . Adam is used for optimization.

4 Experiments

In this Section, we first describe the experimental conditions (4.1), then give the results and analyses of the ablation studies (4.2) as well as the empirical study on the length of temporal window (4.3), then compare our method with related works (4.4). Notably, due to the page limits, the details on hyper-parameter selections, the empirical study on different token communication strategies, the discussion on the computational complexities, as well as more visualization comparisons on the tracking results, are attached in the supplementary materials.

4.1 Experimental Conditions

Two video databases, i.e., the 300VW [Shen *et al.*, 2015] database and the TF [FGNET, 2014] database, are used to evaluate performance of the proposed method. The 300VW database, i.e., the most widely-used database in the field of facial landmark tracking, contains 114 “in-the-wild” videos, each frame is annotated with 68 facial landmarks. 50 videos are divided as the training set and 64 videos for testing. The testing set can be further categorized as three subsets, i.e., S1, S2 and S3, according to their challenging levels.

To be consistent with the experimental conditions of previous works [Liu *et al.*, 2018; Yin *et al.*, 2020], for experiments on the 300VW database, both the 300VW training set and the 300W [Sagonas *et al.*, 2016] training set is used for training. The 300W database is an image database with 3,148 training images, without any temporal information. We just make

multiple copies of an image to be a pseudo static video to pre-train our tracker, then fine-tuning it on the 300VW training set. We adopt all of the 68 landmarks when testing on the 300VW database, and adopt 7 landmarks common to the 300VW and the TF databases for testing on the TF database. All faces are cropped from the facial bounding boxes and resized to 256×256 pixels, then fed into the network for training. Details on hyper-parameter selection are attached in supplementary materials.

We evaluate the tracking performance on both accuracy and stability. Accuracy reflects the closeness of the predicted landmark coordinates to the ground truths. We use Normalized Root Mean Squared Error (NRMSE) [Sagonas *et al.*, 2016] as the accuracy metric. A lower value of NRMSE corresponds to a better accuracy performance. Stability reflects the consistency of movement between predicted landmarks and ground truths. The stability error [Tai *et al.*, 2019] is defined as the error of landmark displacement between the tracking results and the ground truths. A lower value of the stability error corresponds to a better stability performance.

Settings	300VW S1		300VW S2		300VW S3		TF	
	N	S	N	S	N	S	N	S
<i>BL</i>	3.31	0.90	3.45	0.93	4.42	1.85	2.02	0.66
<i>BL+TE^{-a}</i>	3.26	0.87	3.40	0.91	4.39	1.84	2.02	0.65
<i>BL+TE^{-r}</i>	3.06	0.80	3.01	0.76	4.10	1.61	1.98	0.57
<i>BL+TE^{-c}</i>	3.10	0.80	3.03	0.77	4.14	1.65	2.00	0.61
<i>BL+TE</i>	2.91	0.74	2.93	0.73	3.98	1.52	1.92	0.47
<i>BL+IT</i>	3.13	0.84	3.06	0.79	4.20	1.70	2.00	0.62
<i>BL+IR</i>	3.16	0.85	3.12	0.83	4.25	1.73	2.01	0.64
<i>BL+IR+IT</i>	3.09	0.82	2.99	0.76	4.12	1.64	1.98	0.58
<i>BL+TE+IR</i>	2.82	0.72	2.85	0.70	3.93	1.47	1.90	0.41
<i>BL+TE+IT</i>	2.80	0.72	2.81	0.69	3.88	1.44	1.88	0.38
<i>BL+TE+IR+IT</i>	2.74	0.70	2.74	0.67	3.80	1.42	1.86	0.35

Table 1: NRMSE (N/%) and stability error (S/%) of each experimental group.

4.2 Ablation Study

We conduct multiple groups of experiments for ablation studies. The NRMSE and the stability errors of these experiments are recorded in Table 1. The experimental settings are as the following: *BL* denotes the baseline method, which only preserves the backbone 1D representation encoder as well as the 1D heatmap decoders. This makes a simple facial landmark detection method without any temporal modeling and structural modeling modules; *TE* denotes a complete preservation of the temporal modeling mechanisms of 1DFormer; *TE^{-a}*, *TE^{-r}*, and *TE^{-c}* are three ablation terms from *TE*: *-a* replaces the attention mechanism with a simple token mixing communication mechanism, i.e., directly adding the feature from each time step with the average feature of the time window; *-r* removes the recurrently token mixing strategy and simply takes the output features of the backbone encoder from one temporal window as the inputs for attention mechanisms, instead of delivering long-term historical information recurrently as Eq. (2) does; *-c* removes confidence



Figure 4: Visualization of the tracking results on a challenging video clip from the 300VW S3. The red points and green points are respectively the tracking results without / with the recurrently token mixing strategy.

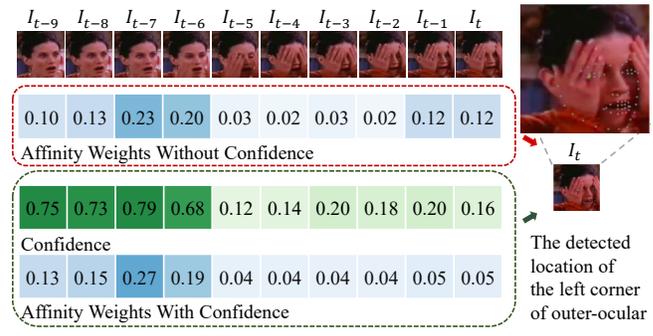


Figure 5: Visualization of the attention weights as well as the tracked results for an occluded landmark, i.e., the left corner of outer-ocular, from a challenging movie clip. The red point and green point are respectively the tracking results without / with the help of the confidence branch.

branches from our attention mechanism in both training and testing phases. *IR* and *IT* respectively denote the proposed intra-group and inter-group geometric encoding mechanisms. Here, *BL+TE+IR+IT* corresponds to the full method of our work. To offset the randomness in the experiment, we conducted 10 independent repeated experiments and reported the average value for each metric in Table 1. From Table 1, we have the following observations and analyses:

First, comparing between the results of *BL+TE* and *BL*, we find that the temporal modeling mechanism *TE* can significantly improve the accuracy and stability performance for facial landmark tracking. This results demonstrate that the proposed temporal modeling mechanisms can effectively integrate the complementary information among multiple video frames to enhance the tracking performance under “in-the-wild” scenarios.

Second, the tracking performance of *BL+TE* shows sig-



Figure 6: Visualization of the tracking results on two challenging video frames, the former from a movie while the latter from the 30VW S3. The red points and green points are respectively the results of the tracker without / with structural modeling mechanisms.

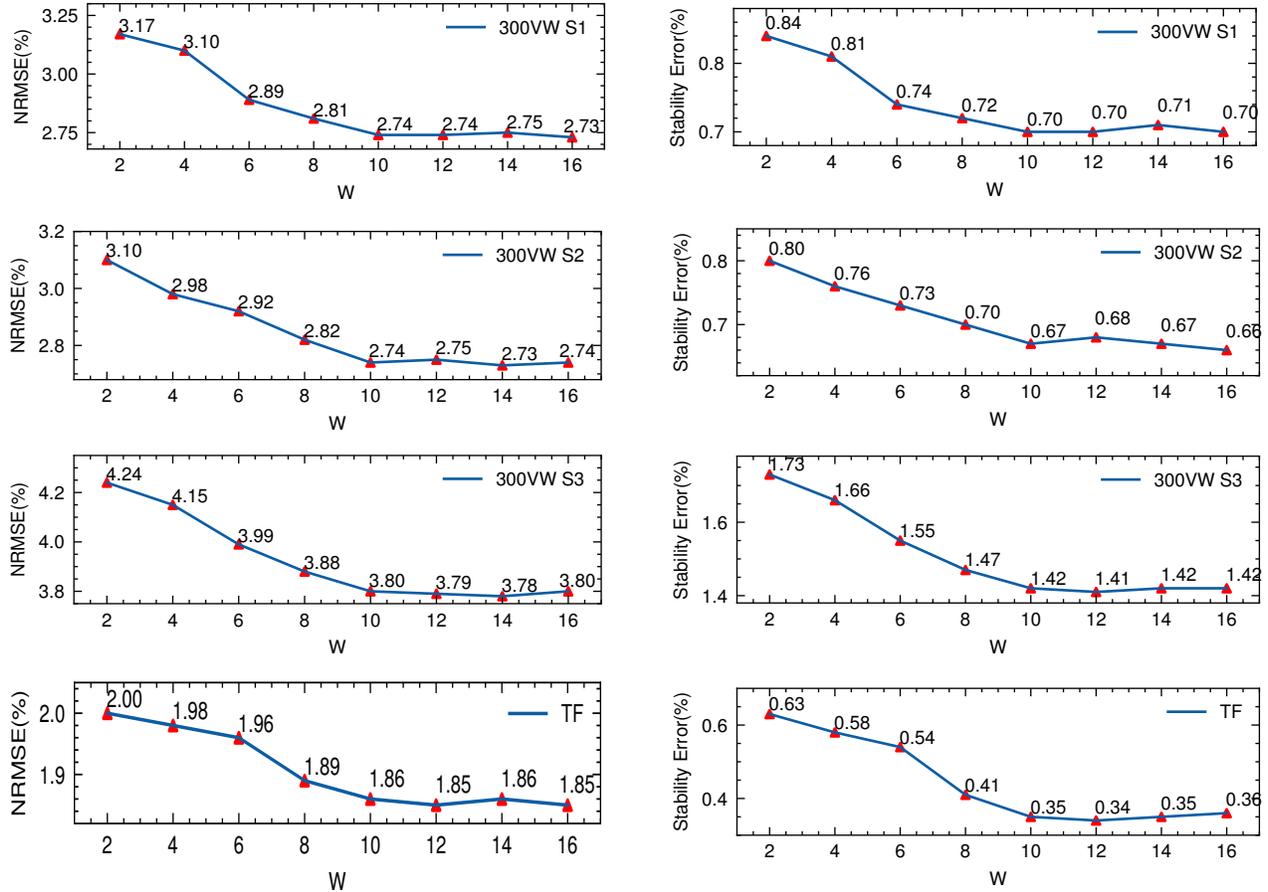


Figure 7: NRMSE (%) and stability errors (%) on 300VW S1, S2, S3, and the TF database with different length (W) of the temporal window.

nificant superiority than $BL + TE^{-a}$, demonstrating the effectiveness of the attention mechanism. This is because the attention mechanism can dynamically adjust the weights of feature fusion, enabling it to adapt to a wide range of imaging conditions and facial geometric variations.

Third, we can observe that recurrently token mixing strategy contributes to the tracking performance by the result that $BL+TE$ outperforms $BL+TE^{-r}$. In the example of Figure 4, without the proposed strategy, the tracking performance at the t th frame is very poor since all of the frames in the temporal window, i.e., from $t - 9$ to t for a window length of 10, suffer from severe occlusions; in contrast, if we add this strategy, we can deliver informative features out of the current window. In this case, information delivered from historical frames, e.g., the $t - 12$ th frame without heavy occlusions, may help to improve the robustness of the tracking performance.

Fourth, the contribution of the confidence branch to the attention mechanism is demonstrated by the comparison between $BL+TE$ and $BL+TE^{-c}$ in Table 1. Figure 5 visualizes the attention weights given by $BL+TE$ and $BL+TE^{-c}$ for an occluded facial landmark, namely the left corner of outer-ocular in this case, from a challenging movie clip. We can find that the confidence branch helps the attention mech-

anism to assign high attention weights if the landmark is visible and assign low weights if it is occluded, providing a plausible guidance to fuse features from different time steps, and thus promoting the localization performance for the occluded frames.

Fifth, we find that both of the intra-group (IR) and inter-group (IE) geometric encoding mechanisms promote the tracking performance. Combined with the illustration in Figure 6, we can find significant improvement brought by our structure modeling modules. The reason lies that, the IR captures the structure patterns of a local component while the IE models the global face patterns, enhancing robustness of 1D landmark representations to challenging imaging conditions.

4.3 Empirical Study on the Length of Temporal Window

We make an empirical study on the length of the temporal window (W) in the recurrently token mixing strategy. From Figure 7, we have the following observations. First, when increasing W from 2 to 10, the tracking accuracy and stability boost significantly. The reason lies that, when W is very small, increasing the value of W can provide necessary temporal information for the Transformer to capture. Second, the tracking performance quickly goes to a converge with the in-

creasing of W . For example, when W is 10, the results have no significant gap with $W = 16$. The reason may be that, as W reaches an appropriate value, e.g., 10 in our experiment, the recurrent information delivery strategy from window to window can play an effective role in capturing the long-term temporal patterns of landmark sequences. Thus, there is no need to continuously increase the value of W and a moderate value of W may also bring a good tracking performance comparable to a large W . This is an advantage which allows for an acceptable computational complexity of our method.

4.4 Comparison with Related Works

We compare our method with the state-of-the-art methods, which include coordinate regression methods, e.g., TSCN [Simonyan and Zisserman, 2014], TSTN [Liu *et al.*, 2018], GAN [Yin *et al.*, 2019], and MSKI [Zhu *et al.*, 2022b]; 2D heatmap regression methods, e.g., FHR [Tai *et al.*, 2019], FHR+STA [Tai *et al.*, 2019], ADC [Chandran *et al.*, 2020], SCPAN [Wan *et al.*, 2021], and SAAT [Zhu *et al.*, 2021]; 1D heatmap regression methods, e.g., the Tracker based on Attentive One-dimensional Heatmap Regression (T-AOHR) [Yin *et al.*, 2020] and SimCC [Li *et al.*, 2022c]; as well as the hybrid method of 2D and 1D heatmap regression, i.e., HybridMatch (HM) [Kang *et al.*, 2023]. From these methods, FHR, ADC, SCPAN, SAAT, SimCC, MSKI, and HM are detection methods, which only consider spatial modeling from a static image or video frame; TSCN, TSTN, FHR+STA, GAN, and T-AOHR are tracking methods which consider both spatial and temporal modeling of a video clip.

Method	year	300VW S1		300VW S2		300VW S3		TF	
		N	S	N	S	N	S	N	S
TSCN	2014	12.54	-	7.25	-	13.13	-	-	-
TSTN	2018	5.36	-	4.51	-	12.84	-	2.13	-
FHR	2019	4.82	2.67	4.23	1.77	7.09	4.43	2.07	0.97
FHR+STA	2019	4.21	1.58	4.02	1.09	5.64	2.62	2.10	0.69
GAN	2019	3.50	0.89	3.67	0.84	4.43	1.82	2.03	0.59
T-AOHR	2020	3.06	0.84	3.17	0.87	4.12	1.78	1.97	0.64
SCPAN	2021	4.49	-	4.23	-	5.87	-	-	-
SAAT	2021	3.46	-	3.41	-	5.23	-	-	-
SimCC	2022	4.07	1.22	4.12	1.09	5.13	2.01	2.09	0.73
MSKI	2022	3.89	1.49	3.94	1.23	5.07	2.07	2.19	-
HM	2023	3.15	0.89	3.23	0.92	4.28	1.80	1.99	0.67
Ours	2024	2.74	0.70	2.74	0.67	3.80	1.42	1.86	0.35

Table 2: NRSME (N %) and stability error (S %) of the proposed tracker and the compared methods on the 300VW and the TF databases.

Table 2 lists the NRMSE performance and stability errors of the proposed tracker and the compared methods on the 300VW and the TF database, respectively. The evaluation results of the compared methods are directly copied from literatures, except for SimCC and HM. Due to the lack of published results of these two methods on the respective databases, we just re-implement them under the same experimental condition as our method. Experimental results from Table 2 show that our method outperforms the compared methods on both tracking accuracy and stability. Our method outperforms the

NRMSE of T-AOHR, which achieves the best overall performance among the compared methods, by 10.46%, 13.56%, 7.77%, and 5.58% on the 300VW scenarios 1, 2, 3, and TF databases, respectively; we also outperforms T-AOHR at the stability performance by 16.67%, 22.99%, 20.22%, and 45.31% on the respective databases. Although T-AOHR applied 1D representations on the facial landmark tracking task and achieved remarkable performance, it ignored a deep exploration on temporal and structural modeling for multiple landmarks, as analyzed in Section 2. We address these weaknesses by developing a Transformer architecture to release the good potentials of 1D representations on modeling the long-range sequential patterns as well as the local and global geometric patterns of facial landmarks, thus achieving significant performance boosts compared to T-AOHR and other related works. We also make visualization comparisons of our method and the related works by rendering their tracking results, please see supplementary materials for details.

5 Conclusions

On addressing the weaknesses of current 1D heatmap regression methods and fully exploring the good potentials of 1D landmark representations on temporal and structure modeling of multiple facial landmarks, we propose a facial landmark tracking method based on a new Transformer architecture, namely 1DFormer, to model long-range temporal patterns as well as the local and global facial structures via token communications in the temporal and spatial dimensions, respectively. Specifically, a confidence-enhanced multi-head attention mechanism with a recurrently token mixing strategy is proposed for temporal modeling; an intra-group and an inter-group geometric encoding mechanism are presented for structure modeling. Experimental results on the 300VW and the TF databases demonstrate that our method achieves state-of-the-art performance for facial landmark tracking with a good modeling capability on the temporal dynamics as well as the geometric patterns of facial landmarks.

Contribution Statement

Shi Yin and Shijie Huang contributed equally to this work.

Acknowledgements

This work was supported by the National Key R&D Program of China (2022YFB4500600)

References

- [Bulat and Tzimiropoulos, 2017] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *ICCV*, pages 1021–1030, 2017.
- [Chandran *et al.*, 2020] Prashanth Chandran, Derek Bradley, Markus H. Gross, and Thabo Beeler. Attention-driven cropping for very high resolution facial landmark detection. In *CVPR*, pages 5860–5869, 2020.
- [Chi *et al.*, 2022] Cailong Chi, Dong Zhang, Zhesi Zhu, Xingzhi Wang, and Dah-Jye Lee. Human pose estimation

- for low-resolution image using 1-d heatmaps and offset regression. *Multimedia Tools and Applications*, pages 1–19, 2022.
- [Chrysos *et al.*, 2018] Grigorios G Chrysos, Epameinondas Antonakos, Patrick Snape, Akshay Asthana, and Stefanos Zafeiriou. A comprehensive performance evaluation of deformable face tracking “in-the-wild”. *IJCV*, 126(2-4):198–232, 2018.
- [FGNET, 2014] FGNET. Talking face video. http://www-prima.inrialpes.fr/FGnet/data/01-TalkingFace/talking_face.html, 2014.
- [Kang *et al.*, 2023] Seungyoon Kang, Minhyun Lee, Minjae Kim, and Hyunjung Shim. Hybridmatch: Semi-supervised facial landmark detection via hybrid heatmap representations. *IEEE Access*, 11:26125–26135, 2023.
- [Li *et al.*, 2013] Y. Li, S. Wang, Y. Zhao, and Q. Ji. Simultaneous facial feature tracking and facial expression recognition. *TIP*, 22(7):2559–2573, 2013.
- [Li *et al.*, 2021] Jiefeng Li, Siyuan Bian, Ailing Zeng, Can Wang, Bo Pang, Wentao Liu, and Cewu Lu. Human pose regression with residual log-likelihood estimation. In *ICCV*, pages 11005–11014, 2021.
- [Li *et al.*, 2022a] Hui Li, Zidong Guo, Seon-Min Rhee, Seungju Han, and Jae-Joon Han. Towards accurate facial landmark detection via cascaded transformers. In *CVPR*, pages 4166–4175, 2022.
- [Li *et al.*, 2022b] Jinpeng Li, Haibo Jin, Shengcai Liao, Ling Shao, and Pheng-Ann Heng. Repformer: Refinement pyramid transformer for robust facial landmark detection. In *IJCAI*, pages 1088–1094, 2022.
- [Li *et al.*, 2022c] Yanjie Li, Sen Yang, Peidong Liu, Shoukui Zhang, Yunxiao Wang, Zhicheng Wang, Wankou Yang, and Shu-Tao Xia. Simcc: A simple coordinate classification perspective for human pose estimation. In *ECCV*, pages 89–106, 2022.
- [Liu *et al.*, 2018] Hao Liu, Jiwen Lu, Jianjiang Feng, and Jie Zhou. Two-stream transformer networks for video-based face alignment. *TPAMI*, 40(11):2546–2554, 2018.
- [Liu *et al.*, 2022] Shengli Liu, Xiaowen Zhu, Zewei Cao, and Gang Wang. Deep 1d landmark representation learning for space target pose estimation. *Remote Sensing*, 14(16):4035, 2022.
- [Moon and Lee, 2020] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single RGB image. In *ECCV*, Lecture Notes in Computer Science, pages 752–768, 2020.
- [Newell *et al.*, 2016] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, pages 483–499, 2016.
- [Peng *et al.*, 2016] Xi Peng, Rogerio S Feris, Xiaoyu Wang, and Dimitris N Metaxas. A recurrent encoder-decoder network for sequential face alignment. In *ECCV*, pages 38–56, 2016.
- [Sagonas *et al.*, 2016] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: Database and results. *Image and vision computing*, 47:3–18, 2016.
- [Shen *et al.*, 2015] Jie Shen, Stefanos Zafeiriou, Grigoris G Chrysos, Jean Kossaifi, Georgios Tzimiropoulos, and Maja Pantic. The first facial landmark tracking in-the-wild challenge: Benchmark and results. In *ICCV Workshops*, pages 50–58, 2015.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, pages 568–576, 2014.
- [Tai *et al.*, 2019] Ying Tai, Yicong Liang, Xiaoming Liu, Lei Duan, Jilin Li, Chengjie Wang, Feiyue Huang, and Yu Chen. Towards highly accurate and stable face alignment for high-resolution videos. In *AAAI*, pages 8893–8900, 2019.
- [Wan *et al.*, 2021] Jun Wan, Hui Xi, Jie Zhou, Zhihui Lai, Witold Pedrycz, Xu Wang, and Hang Sun. Robust and precise facial landmark detection by self-calibrated pose attention network. *IEEE Transactions on Cybernetics*, pages 1–15, 2021.
- [Wang *et al.*, 2021a] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 43(10):3349–3364, 2021.
- [Wang *et al.*, 2021b] Qingzhong Wang, Pengfei Zhang, Haoyi Xiong, and Jian Zhao. Face. evolve: A high-performance face recognition library. *arXiv preprint arXiv:2107.08621*, 2021.
- [Wu and Ji, 2019] Yue Wu and Qiang Ji. Facial landmark detection: A literature survey. *IJCV*, 127(2):115–142, 2019.
- [Wu *et al.*, 2014] Yue Wu, Ziheng Wang, and Qiang Ji. A hierarchical probabilistic model for facial feature detection. In *CVPR*, pages 1781–1788, 2014.
- [Xiong and De la Torre, 2013] Xuehan Xiong and Fernando De la Torre. Supervised descent method and its applications to face alignment. In *CVPR*, pages 532–539, 2013.
- [Xiong *et al.*, 2020] Yilin Xiong, Zijian Zhou, Yuhao Dou, and Zhizhong Su. Gaussian vector: An efficient solution for facial landmark detection. In *ACCV*, volume 12626, pages 70–87, 2020.
- [Yin *et al.*, 2019] Shi Yin, Shangfei Wang, Guozhu Peng, Xiaoping Chen, and Bowen Pan. Capturing spatial and temporal patterns for facial landmark tracking through adversarial learning. In *IJCAI*, pages 1010–1017, 2019.
- [Yin *et al.*, 2020] Shi Yin, Shangfei Wang, Xiaoping Chen, Enhong Chen, and Cong Liang. Attentive one-dimensional heatmap regression for facial landmark detection and tracking. In *ACM MM*, pages 538–546, 2020.

- [Zhu *et al.*, 2021] Congcong Zhu, Xiaoqiang Li, Jide Li, and Songmin Dai. Improving robustness of facial landmark detection by defending against adversarial attacks. In *ICCV*, pages 11731–11740, 2021.
- [Zhu *et al.*, 2022a] Congcong Zhu, Xiaoqiang Li, Jide Li, Songmin Dai, and Weiqin Tong. Reasoning structural relation for occlusion-robust facial landmark localization. *Pattern Recognition*, 122:108325, 2022.
- [Zhu *et al.*, 2022b] Congcong Zhu, Xiaoqiang Li, Jide Li, Songmin Dai, and Weiqin Tong. Multi-sourced knowledge integration for robust self-supervised facial landmark tracking. *IEEE Transactions on Multimedia*, pages 1–13, 2022.