

RePFormer: Refinement Pyramid Transformer for Robust Facial Landmark Detection

Jinpeng Li¹, Haibo Jin², Shengcai Liao^{3*}, Ling Shao⁴ and Pheng-Ann Heng^{1,5}

¹Department of Computer Science and Engineering,
The Chinese University of Hong Kong, Hong Kong, China

²Department of Computer Science and Engineering,
The Hong Kong University of Science and Technology, Hong Kong, China

³Inception Institute of Artificial Intelligence (IIAI), UAE

⁴Terminus Group, China

⁵Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology,
Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China
jpli21@cse.cuhk.edu.hk, haibo.nick.jin@gmail.com, {scliao, ling.shao}@ieee.org,
pheng@cse.cuhk.edu.hk

Abstract

This paper presents a Refinement Pyramid Transformer (RePFormer) for robust facial landmark detection. Most facial landmark detectors focus on learning representative image features. However, these CNN-based feature representations are not robust enough to handle complex real-world scenarios due to ignoring the internal structure of landmarks, as well as the relations between landmarks and context. In this work, we formulate the facial landmark detection task as refining landmark queries along pyramid memories. Specifically, a pyramid transformer head (PTH) is introduced to build both homologous relations among landmarks and heterologous relations between landmarks and cross-scale contexts. Besides, a dynamic landmark refinement (DLR) module is designed to decompose the landmark regression into an end-to-end refinement procedure, where the dynamically aggregated queries are transformed to residual coordinates predictions. Extensive experimental results on four facial landmark detection benchmarks and their various subsets demonstrate the superior performance and high robustness of our framework.

1 Introduction

Facial landmark detection aims to localize a set of predefined key points on 2D facial images. This task has attracted significant attention due to its wide range of applications [Ko, 2018]. With the remarkable success of convolutional neural networks (CNNs), modern facial landmark detectors [Wang *et al.*, 2019a; Qian *et al.*, 2019; Jin *et al.*, 2021] have achieved encouraging performance in constrained environments. However, as fundamental components in real-world facial applications, facial landmark detectors need to be able to consistently

*Shengcai Liao is the corresponding author.



Figure 1: Facial landmark detection results on WFLW. Yellow arrows point to the wrong results. (a) Ground truth. (b) Results of regression-based method. (c) Results of heatmap-based method. (d) Our results. Methods (b) and (c) ignore the internal structures of landmarks, while our method explores the internal structure of landmarks and builds the relations between landmarks and cross-scale contexts, thereby achieving more robust detection results.

generate robust results for complex situations, which remains a significant challenge for existing methods. Figure 1 illustrates some examples of challenging face images and the detection results of existing algorithms and our method.

According to the ways of generating landmark coordinates from feature maps, previous works can be generally classified into three categories: heatmap-based methods, regression-based methods, and hybrid methods. Heatmap-based methods [Wang *et al.*, 2019a; Dong and Yang, 2019] treat facial landmark detection as a segmentation task, where pixels belong to the classes

of landmarks and background. These methods mainly face two difficulties: 1) To reduce the quantization error in CNN backbones, they usually leverage single-scale high-resolution heatmaps [Ronneberger *et al.*, 2015] to represent the surrogate results of landmarks, which introduces high computational costs and prevents them from fully exploring the pyramid features. 2) They lack global constraints on facial landmarks and ignore the relationships between image contexts and landmarks, which decreases their robustness to complex scenes. Regression-based methods [Lv *et al.*, 2017] directly transform image features into landmark coordinates with high efficiency. However, a single regression step is difficult to achieve satisfactory performance. Cascaded regression steps [Lv *et al.*, 2017; Feng *et al.*, 2018] are leveraged to refine results, but they are not fully end-to-end trainable and their performance is prone to saturation. Recently, Sun *et al.* [Sun *et al.*, 2018] combined heatmap and regression together to inherit the advantages of both. However, their network designs are imbalanced, with most of the computational costs concentrated on the backbone. Tan *et al.* [Tan *et al.*, 2020] demonstrated that imbalanced network architectures can limit the detection performance. Moreover, existing detectors often focus on directly learning the landmark representation from pure image features, and ignore the dynamics and interactivity of landmarks and context information.

To address above limitations, we propose a novel facial landmark detector named the Refinement Pyramid Transformer (RePFormer) which focuses on robust detection for complex scenarios. We treat facial landmarks as attention quires, and formulate the facial landmark detection as a task of refining landmark queries through pyramid memories in an end-to-end trainable procedure. Specifically, we propose a pyramid transformer head (PTH) to mimic the zooming-in mechanism. In PTH, cross-scale attention constructs pyramid memories to combine long-range context and propagate high-level semantic information into lower levels. Then, landmark-to-landmark and landmark-to-memory attentions are employed in each PTH stage to enable dynamic interactions of landmarks and pyramid memories, which helps to regress the landmarks in various scenarios. Besides, we design a dynamic landmark refinement (DLR) method to refine the landmark queries by predicting residual coordinates and dynamically aggregating queries in an end-to-end trainable way. The residual coordinates prediction decomposes the regression into multiple steps, and each step refines the landmark queries based on a specific level in pyramid memories, which fully leverages the multi-level information in pyramid memories and decreases the difficulty of prediction. We conduct extensive experiments on several facial landmark detection benchmarks. Our models achieve state-of-the-art performance and show strong robustness on complex scenarios.

2 Related Work

Facial Landmark Detection. Early facial landmark detection methods mainly focus on deforming a statistical facial landmark model into a 2D image by an optimization pro-

cedure. These methods typically apply different constraints on the statistical models, such as object shapes and texture features [Wu and Ji, 2019]. However, they are not robust under various scenarios in the wild. Recently, deep neural network (DNN) based methods [Wang *et al.*, 2019a; Jin *et al.*, 2021] show promising performance on this task. Their architectures are usually composed of a CNN backbone and a variant of a heatmap-based or/and regression-based detection head. The pixel values of heatmaps represent the likelihood of landmarks existing in the corresponding positions [Wang *et al.*, 2019a; Dong and Yang, 2019]. A drawback of them is that the down-sampling ratio of heatmaps introduces quantization error. Thus, high-resolution networks, such as U-Net [Ronneberger *et al.*, 2015] and HR-Net [Wang *et al.*, 2019a], are commonly applied to address this difficulty. Instead of utilizing heatmaps as surrogate results, regression-based methods [Lv *et al.*, 2017] directly predict facial landmarks by transforming image features into 2D coordinates. Cascaded regression steps [Lv *et al.*, 2017; Feng *et al.*, 2018] are widely applied to further refine the predicted coordinates. They usually follow an iterative cropping-and-regression pipeline which first crops patches from CNN feature maps and then sends them to carefully designed regressors, such as Recurrent Neural Network (RNN) [Trigeorgis *et al.*, 2016] or Graph Convolutional Network (GCN) [Li *et al.*, 2020]. However, extracting patches limits the long-range information exchange and may hinder the end-to-end training due to the non-differentiable cropping operators.

Vision Transformer. Transformer [Vaswani *et al.*, 2017] is first proposed for sequence-to-sequence tasks in natural language processing. Its network architecture simply consists of self-attention based encoders and decoders to model dependencies among any positions in a sequence. Recently, computer vision, from high-level tasks to low-level tasks, has also benefited from its strong representation capability and long-range interactions [Carion *et al.*, 2020]. The Vision Transformer (ViT) [Dosovitskiy *et al.*, 2021] directly flattens image patches and processes their feature and position embeddings using a pure transformer encoder for the image classification task. DETR [Carion *et al.*, 2020] introduces a transformer-based detector to solve the object detection problem as unordered set prediction. Its fully end-to-end design removes the hand-crafted components, such as anchor assignment and duplication reduction, in modern object detectors. Transformers have also been explored for landmark detection tasks. Yang *et al.* [Yang *et al.*, 2020] leverage a transformer as a non-local module to build long-range spatial dependencies in image features for more explainable human pose estimation. In contrast, the transformer in our model explicitly establishes the relations between landmarks to landmarks, and landmarks to image contexts, and gradually refines the landmark queries along the pyramid transformer memories.

3 Method

3.1 Overview

Problem Formulation. Given an input 2D image I , our RePFormer aims at learning a facial landmark detector $f_\theta(I)$ to predict an ordered set $L \in \mathbb{R}^{N \times D} = \{l_1, \dots, l_N\}$ which rep-

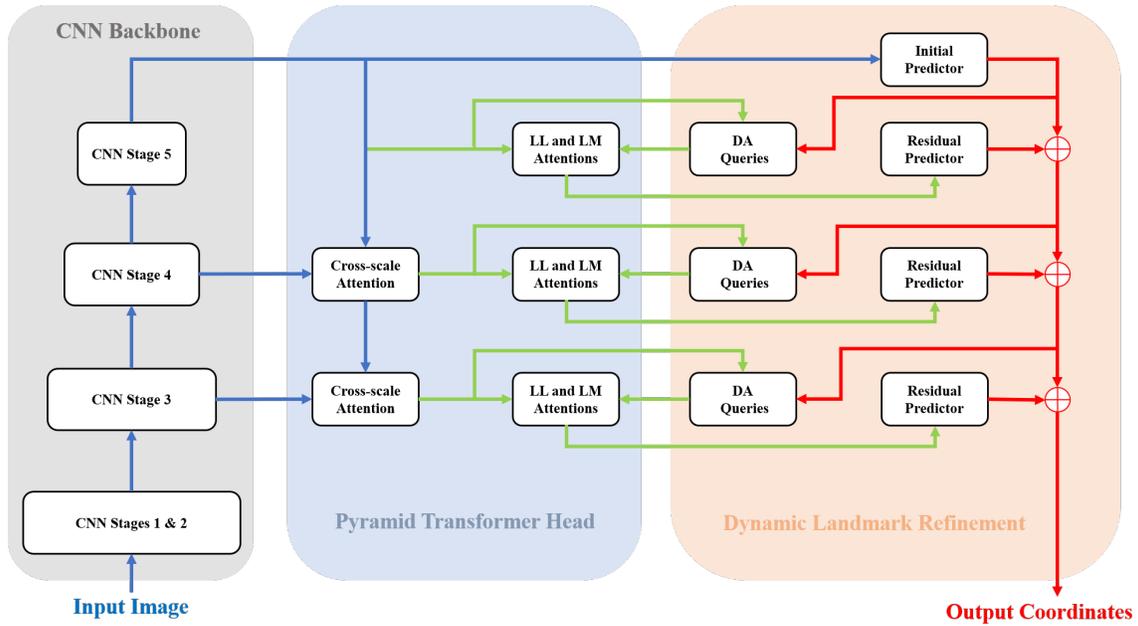


Figure 2: Overview of RePFormer. LL and LM are the landmark-to-landmark attention and landmark-to-memory attention, respectively. DA queries are dynamically aggregated queries. Blue lines represent feature maps. Green lines are the embeddings of memories and landmark queries. Red lines represent coordinates and residual predictions of facial landmarks.

resents N ordered facial landmarks, where each landmark l_i is in D -dimensional coordinates. Note that the transformer-based object detector, DETR [Carion *et al.*, 2020], also formulates general object detection as a set prediction task. However, the predicted set of DETR is unordered and with variable length, since the sequence and number of objects are not fixed in general object detection datasets. In contrast, the facial landmark detection task assumes that each image only contains one main face (if there are several faces in an image, only the largest central one is considered), which is annotated with a fixed number of landmarks. Besides, instead of solving a bipartite matching as in DETR, the label assignment in our framework is simply defined as a fixed one-to-one way.

Overall Architecture. The architecture of our RePFormer is illustrated in Fig. 2. It is built upon a CNN-based feature extractor and a transformer-based detection head. First, the input image is fed into a common CNN backbone which contains several network stages. The feature maps generated by these stages are of various resolutions and semantics, forming hierarchical image features. Next, our PTH module applies cross-scale attentions on these hierarchical feature maps to produce pyramid memories with multi-level semantic information. Then, the PTH stages use landmark-to-landmark and landmark-to-memory attentions to fuse the context information of memories into landmark queries and model the relations among landmarks. After each PTH stage, our DLR module performs mutual updates by predicting the residual coordinates of landmarks based on the status of current queries and evolving the landmark queries based on the coordinates of current landmarks. Thus, the whole facial landmark detection task is formulated as a step-by-step landmark refinement process.

3.2 Pyramid Transformer Head

Cross-Scale Attention. Feature pyramids have been explored for solving various challenges in computer vision, such as detecting objects at various scales [Li *et al.*, 2019] and combining different levels of semantic features for segmentation [Ronneberger *et al.*, 2015]. In this work, our transformer-based detection head, PTH, explicitly employs pyramid features to obtain cross-scale semantic information, mimicking the zooming-in mechanism of human annotators. Our cross-scale attention takes the hierarchical feature set Z as input, and generates the pyramid memory set V following a top-down pathway. Z is obtained by a CNN backbone f_B which is defined as follows:

$$f_B(I) = f_B^M(\dots f_B^i(\dots f_B^1(I))), \quad (1)$$

where M is the number of network stages, and f_B^i is the i th network stage to take the feature maps z_{i-1} as input and generate feature maps z_i as follows:

$$z_i = f_B^i(z_{i-1}). \quad (2)$$

These hierarchical features form the CNN feature set $Z = \{z_M, \dots, z_1\}$, which contains multi-level feature maps with various resolutions and semantics. Top-level features contain strong semantics but less spatial information, making them only suitable for robustly locating the rough positions of landmarks. To better propagate high-level semantics into all levels of features and combine long-range information, our PTH computes the i th memory v_i by applying a cross-scale attention on every two adjacent levels of feature maps z_i and memories v_{i+1} , which is inspired by the top-down Grounding Transformer in [Zhang *et al.*, 2020]. The architecture of cross-scale attention is shown in Fig. 3(a), which takes three

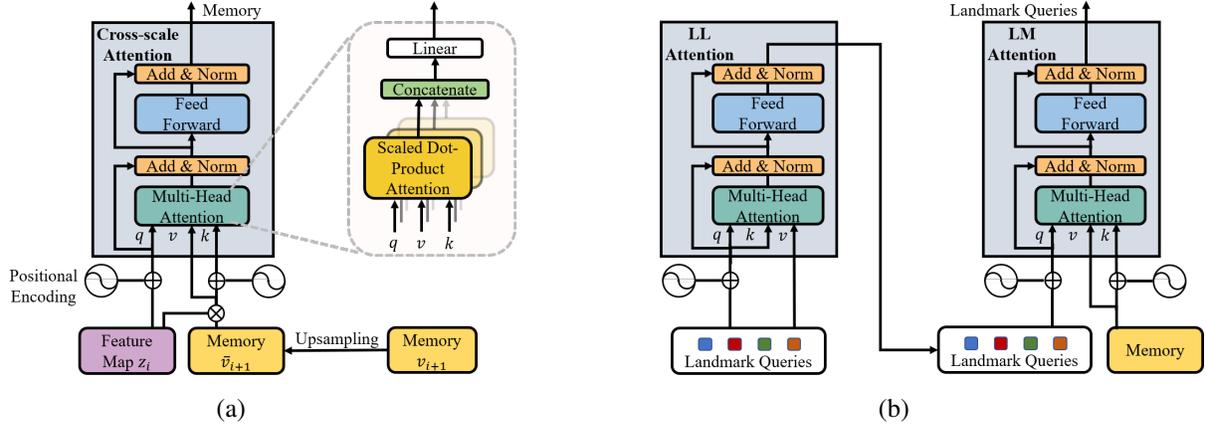


Figure 3: (a) Architecture of Cross-scale attention. (b) Architectures and data flow of LL and LM attentions.

inputs including query q_i , key k_i and value \hat{v}_i . Since the attention operation is permutation-invariant to input elements, q_i is composed of z_i and its fixed positional encodings p_i by $q_i = z_i + p_i$ to maintain relative positions of pixels. k_i and \hat{v}_i are both obtained by fusing the information of z_i and v_{i+1} . A bilinear interpolation upsamples v_{i+1} into \bar{v}_{i+1} with the same resolution as z_i . Then \hat{v}_i is generated by a Hadamard product as $\hat{v}_i = z_i \circ \bar{v}_{i+1}$, where \bar{v}_{i+1} aims to mask output the noisy low-level details outside semantic areas and maintains the precise spatial information within semantic areas. To make the similarity calculation between q_i and k_i semantic and position sensitive, k_i is also supplemented with fixed positional encodings as $k_i = \hat{v}_i + p_i$. In summary, the memory v_i is computed as

$$\hat{z}_i = \text{LN}(z_i + \text{MHA}(q_i, k_i, \hat{v}_i)) \quad (3)$$

$$\begin{cases} v_i = \text{LN}(\hat{z}_i + \text{FFN}(\hat{z}_i)), & \text{if } i < M \\ v_i = z_i, & \text{otherwise,} \end{cases} \quad (4)$$

where MHA is a multi-head attention [Vaswani *et al.*, 2017], FFN is a feed forward network, and LN is a layer normalization. Thus, the pyramid memories form the set $V = \{v_M, \dots, v_1\}$, where each v_i has the same resolution as z_i but with cross-level and long-range information.

PTH Forwarding. Our PTH detection head f_H is composed of cascaded transformer stages, and defined as follows:

$$f_H(V, E^1) = f_H^K(v_{M-K+1}, \dots, f_H^i(v_{M-i+1}, \dots, f_H^1(v_M, E^1))), \quad (5)$$

where K is the number of transformer stages and f_H^i is the i th stage. PTH introduces the landmark query set $E \in \mathbb{R}^{N \times D} = \{e_1, \dots, e_N\}$ to store the states of landmarks, the size of which is the same as the number of landmarks. E can either be initialized from trainable embeddings as in DETR [Carion *et al.*, 2020], or dynamically aggregated from memories as in our method, which will be introduced in next section. PTH takes the pyramid and initial landmark query set E^1 as inputs and repeatedly updates the states of E based on the information from pyramid memories V , in a top-down pathway. This enables it to first robustly locate the rough positions of landmarks in the top-level memory, and then gradually integrates more and more precise spatial information from lower levels

of memories. Specifically, f_H^i is the i th PTH stage to take the memory v_{M-i+1} and the landmark query set E^i in the i th state as inputs and generate the new query state E^{i+1} for the next PTH stage. This is computed by

$$\begin{aligned} E^{i+1} &= f_H^i(v_{M-i+1}, E^i) \\ &= g_{LM}(g_{LL}(E^i), v_{M-i+1}) \end{aligned} \quad (6)$$

where each PTH stage is composed of two successive attentions, i.e. a landmark-to-landmark attention g_{LL} and a landmark-to-memory attention g_{LM} as shown in Fig. 3(b), which is inspired by the transformer decoders in [Vaswani *et al.*, 2017; Carion *et al.*, 2020]. They are both implemented with multi-head attentions. The query, key and value of g_{LL} all come from E^i , and the query, key and value of g_{LM} are the output of g_{LL} , v_{M-i+1} and v_{M-i+1} , respectively. g_{LL} models the dynamic relations among landmark queries using their similarities, and updates their status. g_{LM} computes the relations among landmarks and pyramid memories to endow cross-scale information and long-range image context into landmark queries. These dynamic relations are built on the fly and specific to each image, thus they are more robust in detecting facial landmarks without salient visual features.

3.3 Dynamic Landmark Refinement

Our model directly predicts the coordinates of landmarks, so it can be regarded as a regression-based method. Compared to heatmap-based detectors, a single regression step is hard to generate competitive results. Thus, multiple cropping-then-detection steps [Trigeorgis *et al.*, 2016] are used to improve the performance of regression-based methods in a coarse-to-fine way. However, existing multi-step methods usually have two drawbacks: 1) The hard cropping operator is not differentiable with respect to the input coordinates. 2) The detection is performed only based on the cropped features, without access to the image context information. To address these challenges, our RePformer introduces the DLR module to modulate facial landmark detection into a fully end-to-end refinement process using residual coordinates predictions and dynamically aggregated queries.

Residual Coordinates Prediction. The initial facial landmark set L^0 is predicted based on the highest-level memory v_M . Then, our DLR appends additional predictors to

Method	Backbone	Pose	Expr.	Illu.	WFLW				300W	COFW	AFLW
					M.u.	Occ.	Blur	Full			
DVLN [Wu and Yang, 2017]	VGG-16	11.54	6.78	5.73	5.98	7.33	6.88	6.08	4.66	-	-
LAB [Wu <i>et al.</i> , 2018]	ResNet-18	10.24	5.51	5.23	5.15	6.79	6.32	5.27	3.49	5.58	1.85
Wing [Feng <i>et al.</i> , 2018]	ResNet-50	8.43	5.21	4.88	5.26	6.21	5.81	4.99	-	5.07	1.47
DeCaFA [Dapogny <i>et al.</i> , 2019]	Cascaded U-net	-	-	-	-	-	-	5.01	3.69	-	-
HRNet [Wang <i>et al.</i> , 2019a]	HRNetV2-W18	7.94	4.85	4.55	4.29	5.44	5.42	4.60	3.32	3.45	1.57
AVS [Qian <i>et al.</i> , 2019]	ResNet-18	9.10	5.83	4.93	5.47	6.26	5.86	5.25	4.54	-	-
AVS w/ LAB [Qian <i>et al.</i> , 2019]	Hourglass	8.21	5.14	4.51	5.00	5.76	5.43	4.76	4.83	-	-
AVS w/ SAN [Qian <i>et al.</i> , 2019]	-	8.42	4.68	4.24	4.37	5.60	4.86	4.39	3.86	-	-
AWing [Wang <i>et al.</i> , 2019b]	Hourglass	7.38	4.58	4.32	4.27	5.19	4.96	4.36	3.07	-	-
STYLE [Qian <i>et al.</i> , 2019]	ResNet-18	8.42	4.68	4.24	4.37	5.60	4.86	4.39	4.54	-	-
LUVLi [Kumar <i>et al.</i> , 2020]	DU-Net	-	-	-	-	-	-	4.37	3.23	-	-
DAG [Li <i>et al.</i> , 2020]	HRNet-W18	7.36	4.49	4.12	4.05	4.98	4.82	4.21	3.04	-	-
PIPNet [Jin <i>et al.</i> , 2021]	ResNet-18	8.02	4.73	4.39	4.38	5.66	5.25	4.57	3.36	3.31	1.48
PIPNet [Jin <i>et al.</i> , 2021]	ResNet-50	7.98	4.54	4.35	4.27	5.65	5.19	4.48	3.24	3.18	1.44
PIPNet [Jin <i>et al.</i> , 2021]	ResNet-101	7.51	4.44	4.19	4.02	5.36	5.02	4.31	3.19	3.08	1.42
RePFormer (ours)	ResNet-18	7.38	4.28	4.06	4.04	5.17	4.86	4.20	3.07	3.07	1.44
RePFormer (ours)	ResNet-50	7.31	4.25	4.09	3.94	5.15	4.82	4.14	3.03	3.01	1.43
RePFormer (ours)	ResNet-101	7.25	4.22	4.04	3.91	5.11	4.76	4.11	3.01	3.02	1.43

Table 1: Benchmarking results of state-of-the-art methods and our models on the WFLW including the full set and six subsets, 300W, COFW and AFLW datasets. The best and second best results are marked in colors of red and blue, respectively.

higher PTH stages, and the i th predictor f_P^i only needs to predict the residual coordinates of landmarks $U^i \in \mathbb{R}^{N \times D} = \{u_1^i, \dots, u_N^i\}$ with respect to L^{i-1} , as follows:

$$u_j^i = f_P^i(e_j^i), \tag{7}$$

where u_j^i are the residual coordinates of the j th landmark in the i th stage, and e_j^i is the j th landmark query in the i th stage. The predictor f_P^i is a two-layer FFN which is agnostic to different landmarks, because it only needs to predict the residual values instead of the absolute coordinates. And the i th landmark set L^i is computed by:

$$L^i = \{l_1^{i-1} + u_1^i, \dots, l_N^{i-1} + u_N^i\}. \tag{8}$$

L1 loss is used as the loss function between the ground-truths and L^i generated by the i th PTH stage. Thus, with this step-by-step refinement, our DLR gradually pushes the regressed landmarks closer to the ground-truth coordinates.

Dynamically Aggregated Queries. To adapt the transformer architecture to the residual prediction task, the landmark query set E^i , as the inputs of the i th PTH stage, need to represent the status of the current results L^{i-1} using their semantic and spatial information. To solve this challenge in an end-to-end way, a dynamic aggregation method is proposed to extract landmark queries by aggregating pyramid memories weighted by their relative positional information. This can be seen as a soft version of the "cropping" operator, and the j th query for stage i is computed by as follows

$$e_j^i = \sum_{k \in \Omega} s_{ijk} \cdot v_{M-i+1}^k, \tag{9}$$

where k is a pixel's memory index, Ω is the domain of all indexes, and v_{M-i+1}^k is the k th memory pixel in the stage $M - i + 1$. The normalized similarity s_{ijk} between query e_j^{i-1} and memory pixel v_{M-i+1}^k is represented by

$$s_{ijk} = \frac{\exp(-\|l_j^{i-1} - c_k\| \cdot \tau)}{\sum_{\hat{k} \in \Omega} \exp(-\|l_j^{i-1} - c_{\hat{k}}\| \cdot \tau)}, \tag{10}$$

where $\|\cdot\|$ is the squared L2 norm, and c_k represents the coordinates of memory pixel v_{M-i+1}^k . τ is a temperature parameter which is bigger than 1 to amplify the differences in coordinates. Compared with hard cropping, dynamically aggregated queries are differentiable to coordinates of landmarks, making the entire multi-step refinement procedure fully end-to-end trainable. Besides, these queries are composed of the whole memory to explore both the position sensitive semantic features and global image context information.

4 Experiments

4.1 Implementation Details and Datasets

We use ResNet [He *et al.*, 2016] pretrained on ImageNet [Deng *et al.*, 2009] as the backbone of RePFormer, and the default depth of backbone is 18 unless otherwise specified. The Adam optimizer without weight decay is used to train our models, β_1 and β_2 are set to 0.9 and 0.999, respectively. All models are trained for 360 epochs with a batch size of 16. The initial learning rate is 0.0001, which is decayed by a factor of 10 after 200 epochs. The temperature τ is set to 1000. The L1 loss is used as the loss function for all outputs, and the loss weights are simply set to 1. We conduct experiments on four popular facial landmark detection datasets including WFLW [Wu *et al.*, 2018], 300W [Sagonas *et al.*, 2013], AFLW-Full [Koestinger *et al.*, 2011], and COFW [Burgos-Artizazu *et al.*, 2013]. Most of our settings follow PIPNet [Jin *et al.*, 2021]. All input images are resized to 256x256.

4.2 Comparison with State-of-the-art Approaches

We compare RePFormer with the state-of-the-art facial landmark detection methods using the evaluation metric of normalized mean error (NME).

WFLW. Table 1 shows the performance of state-of-the-art methods and our RePFormer models with three backbones including ResNet-18, ResNet-50 and ResNet-101. Thanks

Modules		300W		
PTH	DLR	Full	Common	Challenge
		3.28	2.92	4.99
✓		3.23	2.88	4.93
	✓	3.22	2.92	4.85
✓	✓	3.07	2.72	4.69

Table 2: Comparative results on 300W by using different components in RePFormer.

to our effective detection head, on the full set of WFLW, our model with the lightweight ResNet-18 backbone already achieves better performance (4.20% NME) than all existing methods including those with much heavier backbones, such as Cascaded U-net and Hourglass. With ResNet-50 and ResNet-101, RePFormer further improves results to 4.14% and 4.11% NME, outperforming the most competitive method, DAG [Li *et al.*, 2020], by 1.7% and 2.4%, respectively. We also conduct experiments on six subsets of WFLW. Our models outperform all previous methods on five subsets, and achieve all the best and second best results, demonstrating the strong performance and robustness of our framework in various evaluation scenarios.

300W. The third-to-last column of Table 1 compares the performance of our models and state-of-the-art methods on the full set of 300W. All the reported results are normalized by the inter-ocular distances. Our RePFormers with ResNet-101 and ResNet-50 backbones achieve the best and second-best results, respectively. Only with the lightweight backbone of ResNet-18, our method already significantly outperforms most of the state-of-the-art methods, such as LAB [Wu *et al.*, 2018], STYLE [Qian *et al.*, 2019], HRNet [Wang *et al.*, 2019a], and PIPNet [Jin *et al.*, 2021]. Note that although RePFormer has a deeper detection head than detectors with specially tailored head for time efficiency, we argue that their architectures of large backbones with small detection heads are not optimal, and our RePFormer head can achieve better performance gain than using heavier backbone. Tan *et al.* [Tan *et al.*, 2020] also showed that balanced architectures can yield higher accuracy-speed ratio. For example, compared to the competitive method, PIPNet [Jin *et al.*, 2021], our RePFormer with ResNet-18 outperforms PIPNet with ResNet-101 by 3.8% with similar inference speed (56 FPS vs. 59 FPS).

COFW. We compare our methods with the state-of-the-art works on the COFW dataset under the intra-data setting which is shown in the second-to-last column of Table 1. The inter-ocular distances are used to normalize the results. Our models with three different backbones achieve the top-3 performances among all methods.

AFLW. Comparative results of our models and state-of-the-art results are shown in the last column of Table 1. Following previous works, all landmark coordinates are normalized by the image size. As can be seen that our RePFormers with different backbones achieve the second-best and third-best results, while only PIPNet [Jin *et al.*, 2021] with ResNet-101 slightly outperforms us (1.42% NME vs. 1.43% NME).

τ	10	100	1000	10000
NME(%)	3.22	3.16	3.07	3.16

Table 3: Performance of different τ on the full set of 300W.

4.3 Ablation Studies

RePFormer Components. Table 2 shows the performance of RePFormer with different components. First, we construct a strong baseline model by appending a three-stage transformer detection head on a CNN backbone. Note that only the feature maps from the last stage of the CNN are used as the memory for the transformer, and landmark embeddings are fixed after training. Thanks to the long-range information modeling of transformer, our baseline model already achieves competitive performance compared to the state-of-the-art methods. Secondly, we apply PTH module on the baseline model. With the help of PTH, the performance of our model is improved on all sets, which demonstrates the effectiveness of pyramid memories with cross-scale information. Thirdly, we integrate DLR into our baseline model. From Table 2, we observe that DLR significantly improves the performance on the challenge subset. Finally, as shown in the last row of Table 2, our RePFormer with PTH and DLR modules achieves large performance improvements over the baseline on all subsets. Note that, compared to the baseline, the performance gain of RePFormer is even larger than the sum of improvements of all individual components, which demonstrates that the two proposed components are complementary to each other and can work together to promote more accurate regression.

Temperature Values. We evaluate the performance of RePFormer with different temperature values on the full set of 300W. τ controls the weight map for dynamically aggregated queries. As shown in Table 3, the performance is continuously improved as the τ value increases from 10 to 1000, but an excessively large τ value decreases the performance. The results indicate that dynamic landmark queries mainly aggregate information from nearby pixels, while long-range information is also valuable for generating accurate results.

5 Conclusion

In this paper, we present a refinement pyramid transformer, RePFormer, for facial landmark detection. The proposed PTH utilizes a cross-scale attention to generate pyramid memories containing multi-level semantic and spatial information. Landmark-to-landmark and landmark-to-memory attentions are employed in each PTH stage to fuse memory information into landmark queries and model long-range relationships between landmarks. A DLR module is introduced to solve the task of landmark regression by a fully end-to-end multi-step refinement procedure where the landmark queries are gradually refined by a dynamic aggregation method through top-down pyramid memories.

Acknowledgments

This work is supported by Hong Kong Research Grants Council under General Research Fund Project No. 14201620.

References

- [Burgos-Artizzu *et al.*, 2013] Xavier P. Burgos-Artizzu, Pietro Perona, and Piotr Dollár. Robust face landmark estimation under occlusion. In *ICCV*, 2013.
- [Carion *et al.*, 2020] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [Dapogny *et al.*, 2019] Arnaud Dapogny, Kevin Bailly, and Matthieu Cord. Decafa: Deep convolutional cascade for face alignment in the wild. In *ICCV*, 2019.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [Dong and Yang, 2019] Xuanyi Dong and Yi Yang. Teacher supervises students how to learn from partially labeled images for facial landmark detection. In *ICCV*, 2019.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [Feng *et al.*, 2018] Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, and Xiao-Jun Wu. Wing loss for robust facial landmark localisation with convolutional neural networks. In *CVPR*, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Jin *et al.*, 2021] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *Int. J. Comput. Vis.*, 2021.
- [Ko, 2018] ByoungChul Ko. A brief review of facial emotion recognition based on visual information. *Sensors*, 2018.
- [Koestinger *et al.*, 2011] Martin Koestinger, Paul Wohlhart, Peter M. Roth, and Horst Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *Workshop on BeFIT*, 2011.
- [Kumar *et al.*, 2020] Abhinav Kumar, Tim K. Marks, Wenxuan Mou, Ye Wang, Michael Jones, Anoop Cherian, Toshiaki Koike-Akino, Xiaoming Liu, and Chen Feng. Luvli face alignment: Estimating landmarks' location, uncertainty, and visibility likelihood. In *CVPR*, 2020.
- [Li *et al.*, 2019] Xiaohan Li, Taotao Lai, Shuaiyu Wang, Quan Chen, Changcai Yang, Riqing Chen, Jinxun Lin, and Fu Zheng. Weighted feature pyramid networks for object detection. In *ISPA/BDCloud/SocialCom/SustainCom*, 2019.
- [Li *et al.*, 2020] Weijian Li, Yuhang Lu, Kang Zheng, Haofu Liao, Chihung Lin, Jiebo Luo, Chi-Tung Cheng, Jing Xiao, Le Lu, Chang-Fu Kuo, and Shun Miao. Structured landmark detection via topology-adapting deep graph learning. In *ECCV*, 2020.
- [Lv *et al.*, 2017] Jiangjing Lv, Xiaohu Shao, Junliang Xing, Cheng Cheng, and Xi Zhou. A deep regression architecture with two-stage re-initialization for high performance facial landmark detection. In *CVPR*, 2017.
- [Qian *et al.*, 2019] Shengju Qian, Keqiang Sun, Wayne Wu, Chen Qian, and Jiaya Jia. Aggregation via separation: Boosting facial landmark detector with semi-supervised style translation. In *ICCV*, 2019.
- [Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [Sagonas *et al.*, 2013] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCV Workshops*, 2013.
- [Sun *et al.*, 2018] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018.
- [Tan *et al.*, 2020] Mingxing Tan, Ruoming Pang, and Quoc V. Le. Efficientdet: Scalable and efficient object detection. In *CVPR*, 2020.
- [Trigeorgis *et al.*, 2016] George Trigeorgis, Patrick Snape, Mihalis A. Nicolaou, Epameinondas Antonakos, and Stefanos Zafeiriou. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *CVPR*, 2016.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
- [Wang *et al.*, 2019a] 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. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [Wang *et al.*, 2019b] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. In *ICCV*, 2019.
- [Wu and Ji, 2019] Yue Wu and Qiang Ji. Facial landmark detection: A literature survey. *Int. J. Comput. Vis.*, 2019.
- [Wu and Yang, 2017] Wenyan Wu and Shuo Yang. Leveraging intra and inter-dataset variations for robust face alignment. In *CVPRW*, 2017.
- [Wu *et al.*, 2018] Wenyan (Wayne) Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *CVPR*, 2018.
- [Yang *et al.*, 2020] Sen Yang, Zhibin Quan, Mu Nie, and Wankou Yang. Transpose: Towards explainable human pose estimation by transformer. *CoRR*, 2020.
- [Zhang *et al.*, 2020] Dong Zhang, Hanwang Zhang, Jinhui Tang, Meng Wang, Xiansheng Hua, and Qianru Sun. Feature pyramid transformer. In *ECCV*, 2020.