

POPoS: Improving Efficient and Robust Facial Landmark Detection with Parallel Optimal Position Search

Chong-Yang Xiang^{1*}, Jun-Yan He², Zhi-Qi Cheng^{3†}, Xiao Wu¹, Xian-Sheng Hua⁴

¹School of Computing and Artificial Intelligence, Southwest Jiaotong University, Chengdu, China

²Institute for Intelligent Computing, Alibaba Group, Shenzhen, China

³School of Engineering and Technology, University of Washington, Tacoma, WA, USA

⁴College of Computer Science and Technology, Zhejiang University, Hangzhou, China

Abstract

Achieving a balance between accuracy and efficiency is a critical challenge in facial landmark detection (FLD). This paper introduces Parallel Optimal Position Search (POPoS), a high-precision encoding-decoding framework designed to address the limitations of traditional FLD methods. POPoS employs three key contributions: (1) Pseudo-range multilateration is utilized to correct heatmap errors, improving landmark localization accuracy. By integrating multiple anchor points, it reduces the impact of individual heatmap inaccuracies, leading to robust overall positioning. (2) To enhance the pseudo-range accuracy of selected anchor points, a new loss function, named multilateration anchor loss, is proposed. This loss function enhances the accuracy of the distance map, mitigates the risk of local optima, and ensures optimal solutions. (3) A single-step parallel computation algorithm is introduced, boosting computational efficiency and reducing processing time. Extensive evaluations across five benchmark datasets demonstrate that POPoS consistently outperforms existing methods, particularly excelling in low-resolution heatmaps scenarios with minimal computational overhead. These advantages make POPoS as a highly efficient and accurate tool for FLD, with broad applicability in real-world scenarios.

Code — <https://github.com/teslatasy/POPoS>

1 Introduction

Facial Landmark Detection (FLD) is a fundamental task in computer vision, essential for identifying and localizing key facial features such as the eyes, nose, and mouth in digital images (Wu and Ji 2018). This crucial process underpins a wide array of advanced applications, including facial recognition systems (Khabarлак and Koriashkina 2021) and three-dimensional face modeling (Zeng et al. 2023). Beyond mere detection, FLD serves as a cornerstone for interpreting human facial expressions (Cheng et al. 2024b,c), analyzing emotional states (Cheng et al. 2024a), and facilitating interactions within digital environments (Cheng et al. 2016,

*Partially completed during a visit to Alibaba and CMU.

†Corresponding author: Zhi-Qi Cheng (zhiqi@uw.edu). Partially conducted as a Project Scientist at CMU. Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

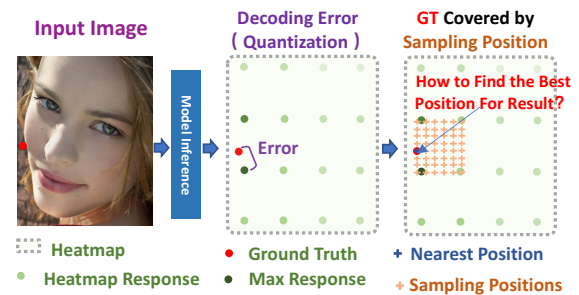


Figure 1: Quantization error in the coordinate decoding process. The discrepancy between the ground truth (GT) coordinates and the highest probability point in the heatmap illustrates the challenge of achieving precise localization and lossless decoding. [Best viewed in color with zoom]

2017a,b,c). Consequently, FLD drives innovation across diverse fields such as human-computer interaction (Li et al. 2024; He et al. 2023, 2024, 2023) and augmented reality systems (Xu et al. 2024a; Tu et al. 2024a,b; Xu et al. 2024b).

The advent of deep learning has advanced FLD research, primarily through two dominant approaches: heatmap-based methods and coordinate regression techniques. Heatmap regression methods (Chandran et al. 2020; Zhang et al. 2020; Jin, Liao, and Shao 2021) have garnered considerable attention due to their ability to generate probabilistic heatmaps for individual landmarks. These methods effectively preserve spatial and contextual information while mitigating the risk of overfitting during training. Additionally, they adeptly capture the inherent uncertainty in landmark positions, enabling robust performance across a wide range of facial expressions and orientations.

Despite their strengths, heatmap-based approaches encounter significant challenges. A primary limitation is quantization error, which becomes pronounced when downscaling heatmaps from their original resolution (Bao et al. 2023). This issue stems from the discretization of continuous spatial information into pixelated representations, leading to potential precision loss during landmark localization. Although maintaining high-resolution heatmaps can mitigate this problem, it imposes substantial computational demands,

rendering it impractical for real-time applications or deployment on resource-constrained devices.

Furthermore, the core challenge in FLD lies in balancing localization accuracy with computational efficiency, particularly in scenarios that require real-time performance or operate under limited processing power. This trade-off is increasingly critical as applications demand both high precision and speed. While pioneering works such as Hourglass (Newell, Yang, and Deng 2016), CrossNet (Zhang et al. 2022), and KeyposS (Bao et al. 2023) have made notable strides in reducing accuracy degradation in low-resolution heatmaps, achieving an optimal balance between accuracy and efficiency remains an open challenge. Current state-of-the-art methods often struggle to maintain high accuracy under low-resolution conditions or stringent computational constraints, indicating a need for further exploration.

This ongoing pursuit for balance in FLD is compounded by the inherent limitations of current encoding and decoding techniques, which further restrict localization performance. Research efforts are actively exploring ways to enhance heatmap-based approaches by leveraging their strengths while mitigating their weaknesses, often through hybrid architectures or advanced post-processing strategies. Fig. 1 illustrates a key challenge in FLD: quantization error in coordinate decoding. The figure highlights the misalignment between the ground truth (GT) landmark coordinates and the heatmap’s highest-probability point. This discrepancy arises from the discretization process, where the continuous spatial domain is mapped onto a discrete pixel grid, compromising localization accuracy. The visualization underscores the importance of advanced decoding strategies to bridge the gap between discrete heatmap representations and the continuous spatial domain of facial features. These limitations emphasize the pressing need for innovative methods to improve accuracy while addressing the challenges of discretization and quantization errors.

To address these interconnected challenges, we present the *Parallel Optimal Position Search (POPoS)* framework, a novel solution designed to resolve the accuracy-efficiency trade-offs inherent in heatmap-based methods. POPoS introduces a set of complementary strategies to tackle the core limitations of traditional FLD approaches. At its foundation, POPoS employs a pseudo-range multilateration technique that leverages the top K points of the heatmap as anchors, effectively reducing the influence of multiple peaks and enhancing landmark localization accuracy. This is further complemented by an anchor-based optimization strategy that integrates global and local refinements, compensating for quantization errors and achieving sub-pixel precision. In addition, POPoS incorporates a single-step parallel computation algorithm, significantly reducing computation time while preserving high accuracy—an approach well-suited to modern GPU architectures. These improvements enable POPoS to excel in low-resolution heatmap scenarios, striking an effective balance between precision and computational efficiency. Consequently, POPoS emerges as a versatile solution applicable to a wide range of real-world settings, from mobile platforms to real-time systems. The key contributions of POPoS are summarized as follows:

- A pseudo-range multilateration approach that improves localization in scenarios with multiple local maxima or ambiguous heatmap regions.
- An anchor-based optimization strategy that synergistically balances global and local optimization, enhancing sub-pixel accuracy.
- A GPU-optimized, single-step parallel computation algorithm that significantly improves efficiency over traditional iterative methods.
- Extensive evaluation across multiple datasets demonstrating superior performance in low-resolution heatmap scenarios, with minimal computational overhead.

2 Related Work

Facial landmark detection research has two key directions: heatmap regression and high-efficiency heatmap decoding.

2.1 Heatmap Regression-Based Methods

Heatmap regression methods (Chandran et al. 2020; Sun et al. 2019; Tang et al. 2018; Zou et al. 2019; Wang, Bo, and Li 2019; Jin, Liao, and Shao 2021; Cheng et al. 2019b,a, 2022; Huang et al. 2020a; Zhang et al. 2022) have progressed by predicting landmarks using high-resolution feature maps. The Stacked Hourglass Network (Newell, Yang, and Deng 2016) and UNet (Ronneberger, Fischer, and Brox 2015) established the foundation by preserving spatial relationships. Subsequent innovations addressed specific challenges: UDP (Huang et al. 2020b) tackled discretization errors, G-RMI (Papandreou et al. 2017) mitigated quantization errors by combining dense heatmaps with keypoint offsets, and the DSNT layer (Nibali et al. 2018) introduced coordinate supervision regularization. PIPNet (Jin, Liao, and Shao 2021) further streamlined the process with simultaneous predictions on low-resolution heatmaps. This progression has consistently enhanced accuracy by leveraging spatial pixel relationships (Zou et al. 2019; Wang, Bo, and Li 2019).

2.2 High Efficiency Heatmap Decoding Methods

While heatmap regression improved prediction accuracy, extracting coordinates from estimated heatmaps remained a critical area for optimization. High-efficiency decoding methods were developed to address this challenge. Darkpose (Zhang et al. 2020) refined heatmap generation and distribution analysis, while FHR (Tai et al. 2019) enhanced precision through improved fractional part estimation. Sub-pixel heatmap regression (Bulat, Sanchez, and Tzimiropoulos 2021) increased accuracy by predicting offsets using soft continuous maxima. KeyPosS (Bao et al. 2023) introduced a novel approach by integrating GPS-inspired multi-point positioning technology for true-range estimation.

Building on these advancements, our Parallel Optimal Position Search (POPoS) method combines pseudo-range multilateration and Potential Position Parallel Sampling. POPoS addresses limitations in predictive accuracy and computational efficiency, particularly in scenarios with low-resolution heatmaps. This innovation not only advances facial landmark detection but also extends its applicability to resource-constrained environments. It marks a step towards

more accessible and efficient high-accuracy detection across diverse applications and hardware configurations.

3 Parallel Optimal Position Search (POPoS)

Despite advancements like HRNet (Sun et al. 2019), heatmap-based facial landmark detection faces persistent challenges. Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, the objective is to locate N_k facial landmarks $\{\beta_k\}_{k=1}^{N_k}$, where $\beta_k = (u_k, v_k)$. Current approaches use CNNs to generate heatmaps $\mathcal{H} \in \mathbb{R}^{h \times w \times N_k}$, with $h = H/\lambda$, $w = W/\lambda$, and λ as the downsampling factor. Key challenges include:

- **Resolution-Efficiency Trade-off.** High-resolution heatmaps ($\lambda \rightarrow 1$) provide precise localization but increase computational load, hindering real-time performance (Newell, Yang, and Deng 2016). Low-resolution heatmaps ($\lambda \gg 1$) are efficient but less accurate, impacting applications in resource-constrained scenarios.
- **Encoding-Decoding Dilemma.** Biased encoding ($\beta_{\text{biased}} = \text{Quan}(\beta/\lambda)$) introduces quantization errors, especially at large λ (Zhang et al. 2020). Unbiased encoding ($\beta_{\text{unbiased}} = \beta/\lambda$) maintains accuracy but results in diffuse heatmaps and increased computational cost (Zhang et al. 2020).
- **Decoding Limitations.** One-hot decoding ($\hat{\beta} = \text{argmax}(\mathcal{H})$) is efficient but prone to discretization errors. Distribution-aware methods (Zhang et al. 2020) perform well at high resolutions but struggle with low-resolution or multi-modal heatmaps, particularly in complex scenarios (Chandran et al. 2020).
- **Optimization Constraints.** The commonly used Mean Squared Error (MSE) loss assumes Gaussian error distribution, often inappropriate for facial landmarks (Zhang et al. 2020). It fails to capture spatial relationships between landmarks, leading to anatomically implausible configurations.

These interconnected challenges hinder robust, accurate, and efficient facial landmark detection in a wide range of scenarios. To address these challenges, we propose Parallel Optimal Position Search (POPoS), a framework that fundamentally reimagines heatmap-based facial landmark detection. POPoS comprises three stages, as shown in Fig. 2. First, the model optimization phase involves generating heatmaps and optimizing them, where the multilateration anchor loss is applied to minimize encoding information loss. Second, the heatmap generation and anchor sampling phase include model inference, distance transformation, and anchor point selection. Third, the sample and search phase involves constructing the pseudo-range positioning equation, parallel sampling, and equivalent single-step parallel computation. Details are in the following sections.

3.1 Model & Optimization

Our adopted HRNet (Bao et al. 2023), specially optimized for memory efficiency, is particularly suitable for mobile and edge computing applications. Its optimizations include early downsampling and a consistent application of 3×3 convolutional kernels, ensuring a balance between performance and

computational efficiency. The optimization objective of current Mean Squared Error (MSE) loss is to minimize the average squared difference between the predicted values and the actual values of the entire heatmap. While the predicted values often biased towards outliers at low resolutions after optimization, which will induce encoding errors. To enhance accuracy of the distance map for anchor points during the decoding step, the inconsistent optimization objectives between the decoding and encoding steps need to be solved. The Multilateration Anchor (MA) loss (Eq.1) is proposed to address the optimization challenge by enhancing accuracy within the top K points, where top K points correspond to the K station anchors at the decoding step:

$$\mathcal{L}_{\text{MA}} = \sum_{i=1}^{N_k} \|\mathcal{H}_i \cdot \text{Mask}_i - \hat{\mathcal{H}}_i \cdot \text{Mask}_i\|_2, \quad (1)$$

where \mathcal{H}_i and $\hat{\mathcal{H}}_i$ are the ground truth and predicted heatmaps, and $\text{Mask}_i \in \mathbb{R}^{h \times w}$ indicates a binary mask which assigns the top K response locations with value one, while the others remain zero, i represents the i -th landmark. Further, the model is jointly optimized using MSE loss and MA loss. The loss L in the training stage is $\mathcal{L} = \mathcal{L}_{\text{MSE}} + w \times \mathcal{L}_{\text{MA}}$, where \mathcal{L}_{MSE} is the MSE loss of the whole heatmap, \mathcal{L}_{MA} is the MA loss, and w is the weight parameter of MA loss.

3.2 Heatmap Generation & Anchor Sampling

In this stage, there are three key steps: 1) heatmap generation, 2) distance transformation, and 3) anchor sampling.

First step is heatmap generation. An input image $I \in \mathbb{R}^{H \times W \times 3}$ is processed through a streamlined HRNet variant (Bao et al. 2023) to extract features. The heatmaps $\mathcal{H} \in \mathbb{R}^{h \times w \times N_k}$ is then generated, where $h = H/\lambda$, $w = W/\lambda$, and N_k represents the number of facial landmarks. The unbiased coordinate encoding is employed during the encoding stage. The coordinate heatmap is generated as a 2-dimensional Gaussian distribution/kernel centered at the labeled coordinate.

The second step is distance transformation, which converts the heatmap into a distance map. The distance between the predicted position and each heatmap pixel, $\mathcal{D}(i, j)$, is calculated as $\mathcal{D}(i, j) = \sqrt{-2\sigma^2 \ln(\mathcal{H}(i, j))}$, where σ is the standard deviation of the Gaussian distribution. Because the heatmap retains the initial positional information entirely, thereby the distance heatmap maintains sub-pixel accuracy.

The third step is anchor sampling. To mitigate the prediction errors induced by heatmap generation, different from the triangulation algorithm (Tai et al. 2019), the top K points of the distance map are selected as anchors. Because the selected top K points are consistent with those top K points in the model optimization stage, they promise to minimize information loss. Meanwhile, these selected top K points are utilized to form distance equations in the next stage.

3.3 Parallel Sampling & Computing

This stage involve two key steps: 1) forming the pseudo-range multilateration equation, and 2) potential position parallel sampling and computing as follows:

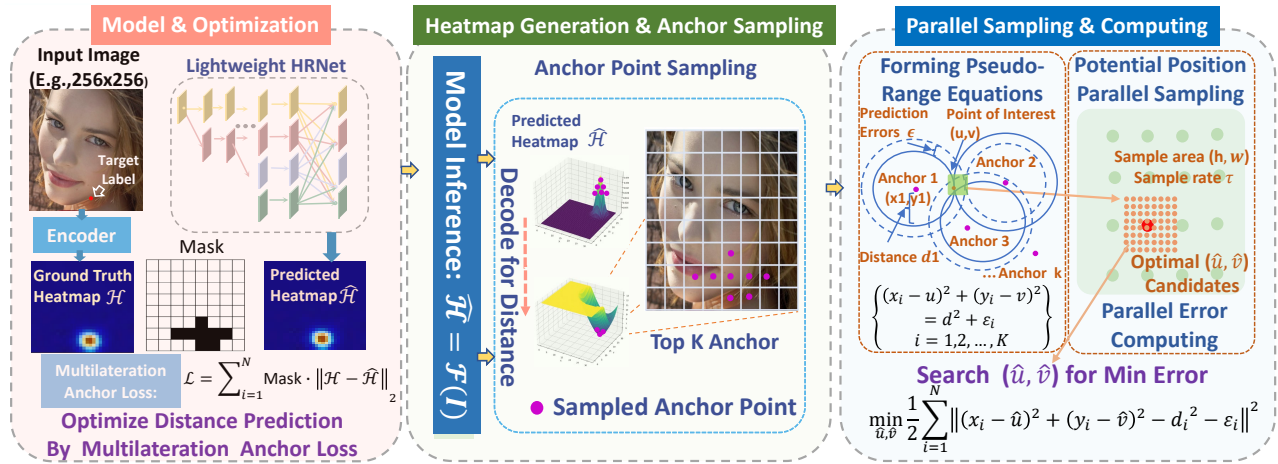


Figure 2: Parallel Optimal Position Search (POPoS) framework. The Multilateration Anchor Loss is designed during the training phase to optimize distance prediction. During model inference, POPoS samples the top-K response positions as anchor points for distance map decoding. [Best viewed in color with zoom.]

Forming Pseudo-range Multilateration Equations The distance map decoding is treated as the pseudo-range multilateration (PRM) problem. Here we utilize selected top K points as anchors, to form a set of distance equations. These distance equations represent the distances between multiple station anchors to the target point. Formally, the top K anchors are $\mathcal{A} = \{A_{(x_0, y_0)}, \dots, A_{(x_K, y_K)}\}$, the real distances between \mathcal{A} and the real position point of interest $\beta(u, v)$ are $\mathcal{D} = \{d_0, \dots, d_K\}$, and the prediction distance error for each anchor is denoted by ε_i . The system of non-linear PRM equations is given by Eq. 2:

$$\{(x_i - u)^2 + (y_i - v)^2 = d_i^2 + \varepsilon_i \mid i = 0, 1, \dots, K\}. \quad (2)$$

Potential Position Parallel Sampling and Computing To solve the pseudo-range multilateration equation and find the point of interest $\hat{\beta} = (\hat{u}, \hat{v})$, the objective function is constructed as Eq. 3:

$$\min_{\hat{u}, \hat{v}} \frac{1}{2} \sum_{i=1}^K \left\| (x_i - \hat{u})^2 + (y_i - \hat{v})^2 - d_i^2 - \varepsilon_i \right\|^2, \quad (3)$$

where ε_i represents the error between the decoded distance and the actual distance. A new method for solving the objective function is proposed, called Potential Position Parallel Sampling and Computing (PPPSC). Through three steps (potential region selection, candidate generation, and Solution of Parallel Computing), it can directly find the point that minimizes the objective function. The detailed steps are outlined and further explained below.

(1) *Potential Region Selection.* First, considering the sub-pixel scenario, there are an infinite number of points existing in the heatmap. Second, as the distance between the sample point with the maximum value point m in the heatmap increases, the probability of sampling point hitting the point of interest decreases. Third, an excess of sampling can lead to redundant computations. Hence, the search area is set as a square region of $h \times w$ centered at point m .

(2) *Candidates Generation.* Within the selected region, sampling is conducted using a mesh with a frequency of τ samples per pixel and obtained a total of N points. Formally, a mesh grid point set $G = \{g_1, \dots, g_N\}$ is generated first, $g_i = (\frac{\Delta \hat{u}_i}{\lambda}, \frac{\Delta \hat{v}_i}{\lambda})$, where $\Delta \hat{u}_i \in [0, \lambda \times h \times \tau]$, $\Delta \hat{v}_i \in [0, \lambda \times w \times \tau]$. And τ is a factor to control the granularity of subpixel sampling. λ is the downsampling factor. Then, the mesh grid point set G is translated to cover the maximum value position (\hat{u}_m, \hat{v}_m) in the grid center, and the candidates $c_i \in C$ is computed as Eq. 4:

$$c_i(\hat{u}_i, \hat{v}_i) = \left(\frac{\Delta \hat{u}_i}{\lambda} + \hat{u}_m - \frac{w}{2}, \frac{\Delta \hat{v}_i}{\lambda} + \hat{v}_m - \frac{h}{2} \right). \quad (4)$$

(3) *Solution of Parallel Computing.* After the sampling, the computing step for one heatmap is formed as Eq. 5. All points from the sampling set C are then inputted into the objective function $\Phi(\mathcal{A}, C)$, which calculates the discrepancy between two distance matrices: one is the Euclidean distance matrix between the sampled points C and the anchor set \mathcal{A} , and the other one is the predicted distance map D of anchor set \mathcal{A} . This results in an error matrix E . The optimal facial landmark position is determined by locating the minimum position index in the matrix E .

$$\begin{aligned} \Phi(\mathcal{A}, C) &= \arg \min_i \sum_k (||C - A||_2 - D)_{ik} \\ &= \min_{\hat{u}, \hat{v}} \sum_{i=1}^K \left\| (x_i - \hat{u})^2 + (y_i - \hat{v})^2 - d_i^2 - \varepsilon_i \right\|^2, \end{aligned} \quad (5)$$

where x_i and y_i are the coordinates of the i -th anchor point. \hat{u} and \hat{v} are the coordinates of the sampled point, d_i is the measured distance from anchor i to the target point, ε_i represents error between the decoded distance and the actual distance for the i -th anchor. The optimization objectives of Eq. 3 and Eq. 5 are consistent, thus PPPSC methods can achieve the optimal approximate solution.

At last, Iterative Gauss–Newton Optimization (IGNO) for PRM is introduced. To validate the effectiveness of our approach, we implemented a general method, IGNO, for solving nonlinear equations of the PRM. Specifically, we utilized the Gauss-Newton iteration method to solve the PRM objective function Eq. 3. To find the point of interest $\hat{\beta} = (\hat{u}, \hat{v})$, the iteration starts at some initial guess $\hat{\beta}_0$, and update $\hat{\beta}_{n+1} = \hat{\beta}_n + \Delta\hat{\beta}$. At each iteration, the update $\Delta\hat{\beta}$ can be calculated from Eq. 6:

$$\left(\sum_{i=1}^k \mathbf{J}_i^{-1} \mathbf{J}_i^T \right) \Delta\hat{\beta} = \sum_{i=1}^K -\mathbf{J}_i^{-1} e_i, \quad (6)$$

where k represents the number of anchors, J is Jacobian matrix of the objective function, e_i represents the objective function value at the i -th iteration. The iteration process halts either when e_i gradually converges or when it reaches the upper limit of iterations. The IGNO methods require computations for each point individually, which has some challenges such as extensive loop operations and GPU acceleration. The proposed PPPSC approach only involves matrix operations, and is thus inherently suitable for parallel computation on GPUs. The computational complexity of the PPPSC approach can be expressed as $4 \times N_a \times N_c$, where N_a represents the number of anchors in \mathcal{A} , and N_c is the number of samples in C . While, IGNO’s complexity, which requires multiplicative operations and matrix-solving computations, is larger than $(24 \times N_a + 30) \times N_{iter}$, where N_{iter} is the iteration count. The PPPSC enables parallel operations on the heatmaps within a batch, and facilitates parallel processing on both the encoding and decoding steps on the GPU. The PPPSC’s complexity is much less than IGNO’s. The following experiment result in Table 2 also prove this conclusion.

4 Experiments

4.1 Dataset and Metric

Datasets Five landmark detection datasets are employed for a comprehensive evaluation of the POPoS framework:

- **COFW** (Artizzu, Perona, and Dollár 2013) comprises 1,345 training and 507 testing images, each with 29 annotated landmarks. This dataset emphasizes scenes with faces under occlusion.
- **AFLW** (Köstinger et al. 2011) includes a diverse set of approximately 25,000 face images from Flickr, annotated with up to 21 landmarks per image, showcasing wide variations in appearance and environmental conditions.
- **COCO-WholeBody** (Jin et al. 2020) offers a collection of over 200,000 labeled images and 250,000 instances across 133 keypoint categories, including comprehensive annotations for face landmarks.
- **300W** (Sagonas et al. 2016) consists of approximately 3,937 face images, each annotated with 68 facial landmarks, encompassing a broad spectrum of identities, expressions, and lighting conditions.
- **WFLW** (Wu et al. 2018) is a widely-used dataset in facial landmark detection, containing 10,000 images (7,500 for training, 2,500 for testing), that each image is annotated with 98 landmarks.

Method	WFLW	300W	AFLW	COFW	Para.	GFlops
TSR (Lv et al. 2017)	-	4.99	2.17	-	-	-
Wing (2018)	-	3.60	1.47	-	91.0M	5.5
ODN (2019)	-	4.17	1.63	5.30	-	-
DeCaFa (2019)	4.62	3.39	-	-	10M	-
DAG (2020)	4.21	3.04	-	4.22	-	-
AWing (2019)	4.36	3.07	1.53	4.94	24.1M	26.7
AVS (2019)	4.39	3.86	1.86	4.43	28.3M	2.4
ADA (2020)	-	3.50	-	-	-	-
LUVLi (2020)	4.37	3.23	1.39	-	-	-
PIPNet-18 (Jin, Liao, and Shao 2021)	4.57	3.36	1.48	-	12.0M	2.4
PIPNet-101 (Jin, Liao, and Shao 2021)	4.31	3.19	1.42	-	45.7M	10.5
DTLD (Li et al. 2022a)	4.08	2.96	1.38	-	-	-
RePFormer (Li et al. 2022b)	4.11	3.01	1.43	-	-	-
SLPT (Xia et al. 2022)	4.14	3.17	-	-	9.98M	-
EF-3-ACR (Fard and Mahoor 2022)	-	3.75	-	3.47	-	-
ADNet-FE5 (Huang et al. 2023)	4.1	2.87	-	-	-	-
ResNet50-FE5 (Huang et al. 2023)	-	4.39	-	-	-	-
HRNet-FE5 (Huang et al. 2023)	-	3.46	-	-	-	-
STAR Loss (Zhou et al. 2023)	4.02	2.87	-	4.62	13.37M	-
RHT-R (Wan et al. 2023)	4.01	3.46	1.99	4.42	-	-
POPos (ours, 32 ²)	4.28	3.38	1.43	3.8	9.7M	1.2
POPos (ours, 64 ²)	3.95*	3.28*	1.34	3.44	9.7M	4.7

Table 1: Comparison with state-of-the-art methods. The results are presented in NME (%). The best results are highlighted in bold. * indicates the result of the IGNO solution.

4.2 Evaluation Metric

Following standard practice in facial landmark detection (Bao et al. 2023), the Normalized Mean Error (NME) is used as evaluation metric, which is defined as $\frac{1}{N_k} \sum_{i=1}^N \frac{\|\beta_i - \hat{\beta}_i\|_2}{d}$, where N_k is the number of landmarks, β and $\hat{\beta}$ represent the ground truth and predicted landmarks, respectively, and d is the normalization factor, typically the interupillary distance.

4.3 Implementation Details

Our proposed POPoS framework builds upon the KeyPosS and MMpose¹ frameworks. Under same condition with KeyPosS’s preprocessing and augmentation techniques, we evaluate the performance at varying heatmap resolutions: 64×64 , 32×32 , 16×16 , 8×8 , and 4×4 pixels. The backbone network is a lightweight variant of HRNet (Sun et al. 2019), which is pre-trained on ImageNet (Russakovsky et al. 2015) and fine-tuned using the Adam optimizer (Kingma and Ba 2015) with a linear-step decay learning rate schedule. The initial learning rate is set at 2.0e-3, and gradually reducing to 1.0e-5 over 100 epochs. The training is performed on a 4 × NVIDIA 3090 GPU server, with a batch size 32. In the decoding stage of POPoS, the key hyperparameters, such as the sampling range ($w=1$) and the number of anchors ($K=10$), are determined through training on the 300W (Sagonas et al. 2016) dataset.

¹<https://github.com/open-mmlab/mmpose>




Figure 3: Keypoint detection accuracy at heatmap resolutions of 64×64 , 32×32 , 16×16 , 8×8 , and 4×4 . Smaller radar maps indicate better performance. [Best viewed in color with zoom]

	FPS	Cost
One-hot	2819	0.36ms
Two-hot (Newell, Yang, and Deng 2016)	1394	0.72ms
DarkPose (Zhang et al. 2020)	186.2	5.3ms
KeyPosS (Bao et al. 2023)	96.5	10.34ms
IGNO for PRM	2.5	405.06ms
POPoS	1301	0.76ms

Table 2: FPS of existing decoding methods and PoPoS.

4.4 Comparison with State-of-the-Art Methods

The effectiveness of the Parallel Optimal Position Search (POPoS) is evaluated against state-of-the-art methods across five pivotal datasets, detailed in Table 1. The NME of POPoS is less than the traditional coordinate regression-based method and other heatmap-based methods. We think that coordinate regression-based methods (TSR, Wing, et al) compromise spatial detail retention, thereby affecting precision in high-resolution scenarios. POPoS’s exceptional performance, particularly on the WFLW, AFLW, COFW, and COCO-Wholebody datasets, emphasizes its improved capability in addressing encoding and decoding errors.

Notably, POPoS achieves impressive performance in WFLW, AFLW, and COFW with fewer parameters and higher computational efficiency. On the 300W dataset, POPoS reaches a competitive result Normalized Mean Error (NME) of 3.28%, highlighting the complexity of the dataset and the challenges with limited training data. Notice that, the NME value of the COCO-Wholebody dataset is only used to analyze the decoding precision of various heatmap resolutions, as seen in Fig. 3.

4.5 Comparison with Decoding Schemes

The key problem addressed in this article is the verification of decoding accuracy at different resolutions. The decoding effects of five different methods were compared across five datasets under various heatmap resolution settings. As shown in Fig. 3, a distinctive advantage of POPoS is found in the integration of the Iterative Gauss–Newton Optimization (IGNO) approximation method with a strategic sampling approach. This integration is especially beneficial in low-resolution heatmaps settings, where high-quality results are delivered efficiently by POPoS.

Additionally, to evaluate the efficiency of our proposed POPoS on heatmap decoding, the processing speed of different decoding methods, with the metric of Frame per sec-




Figure 4: Effect of Gauss-Newton Optimization (IGNO), Multilateration Anchor (MA) loss, and Potential Position Parallel Sampling and Computing (PPS).

ond (FPS), is calculated. Although the current heatmap inference speed reaches 200 FPS (Jin, Liao, and Shao 2021), slow heatmap decoding could potentially become a bottleneck. The decoding speed of five methods, one-hot, two-hot, Dark (Zhang et al. 2020), KeyPosS (Bao et al. 2023), and POPoS are compared in Table 2. The FPS for 256×256 input images with 64×64 heatmaps was measured to assess decoding speed. It is clear that the POPoS is 15 times faster than KeyPos, 6 times faster than Dark (Zhang et al. 2020), and comparable in speed to Two-Hot. Although slower than One-Hot, POPoS has significantly improved accuracy and meets real-time processing requirements.

4.6 Ablation Studies

This section offers an in-depth analysis of the effects of the proposed modules and hyperparameter settings.

Effect of Pseudo-Range Multilateration The effectiveness of Iterative Gauss–Newton Optimization (IGNO) for pseudo-range multilateration (PRM) was evaluated by integrating it into the baseline model, replacing the least squares approximation. The number of iterations was capped at 20 or halted upon convergence. KeyPosS (Bao et al. 2023) served as the baseline. Improvements were observed in the 300W and COFW datasets, as illustrated in Fig. 4. These enhancements are attributed to a reduction in prediction errors achieved through an increased number of anchors.

Effect of MA Loss The impact of Multilateration Anchor (MA) loss on mitigating encoding errors was assessed by re-training the HRNet network across various resolutions. The integration of the multilateration anchor loss resulted in significant improvements in POPoS across all datasets, as depicted in Fig. 4. Specifically, in the 4×4 resolution heatmap, prediction errors decreased by 47% for the 300W dataset and by 58% for the COFW dataset.




Figure 5: Detailed analysis of POPoS performance factors: (a) Effect of sampling area (b) Effect of sampling rate τ (c) Anchor number and (d) MA loss weight, each illustrating their impact on the accuracy of predictions.

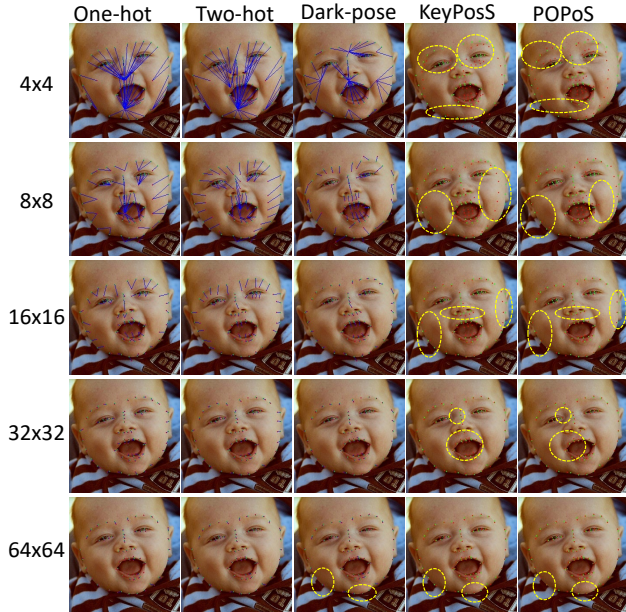


Figure 6: Visualization of POPoS’s performance in low-resolution heatmaps settings: Comparative analysis with one-hot, two-hot, dark-pose, and KeyPosS methods. Ground truth is represented by green dots, predictions by red dots, discrepancies by blue lines, and significant inaccuracies are circled in yellow. [Best viewed in color with zoom]

Effect of PPPSC Potential Position Parallel Sampling and Computing (PPPSC) was evaluated on five datasets at varying resolutions, as presented in Fig. 4. This approach significantly enhanced the stability and performance of the IGN Optimization algorithm. Ablation experiments on the sampling range (h, w) and sampling frequency τ were conducted using the 300W dataset at a resolution of 32×32 . The results, shown in Fig. 5, indicated that the optimal sample area values lie within 1-2 pixels around the maximum value. Increasing the sampling range did not improve decoding accuracy, suggesting that some errors originate in the encoding stage and cannot be corrected by decoding adjustments. Additionally, ablation studies on the sample rate τ demonstrated that optimal performance is achieved with approximately 10 sampling points per pixel length, equivalent to about 1.2 points per original image pixel, as shown in Fig. 5a. Further increases in the number of sampling points

resulted in negligible improvements.

Ablation Study on Hyperparameter Settings The ablation study on the number of anchors N_a in the decoding step, as shown in Fig. 5c, indicates that too few anchors result in inaccurate predictions, while an excessive number of anchors introduces larger calculation errors. For a 32×32 resolution heatmap, the optimal number of anchors was determined to be 25. Additionally, an ablation experiment on the weight of the Multilateration Anchor (MA) loss in the 300W dataset (Fig. 5d) demonstrated consistent improvements in prediction accuracy across various weight ratios, with the most effective ratio being 6. This finding highlights the significant impact of encoding accuracy errors on pseudo-range multilateration and the efficacy of the MA loss in mitigating convergence issues.

4.7 Visualization Analysis

To evaluate POPoS’s adaptability in practical applications, we conducted a comprehensive visualization analysis comparing it with traditional methods such as one-hot, two-hot, DarkPose, and KeyPosS. Focusing on challenging low-resolution heatmaps scenarios depicted in Fig. 6, the analysis revealed that conventional methods struggle with precision at heatmap resolutions of 4×4 and 8×8 . In contrast, POPoS demonstrates superior performance starting from 4×4 and maintains high accuracy at higher resolutions. Additionally, POPoS effectively mitigates localization errors across all tested resolutions, as clearly shown in Fig. 6. These visualizations confirm POPoS’s robustness, highlighting its ability to optimize encoding and decoding processes effectively across diverse application scenarios.

5 Conclusion

We introduce Parallel Optimal Position Search (POPoS), a heatmap-based framework that effectively balances accuracy and efficiency in facial landmark detection (FLD). POPoS employs the *Multilateration Anchor Loss* during encoding to reduce heatmap prediction errors by enhancing anchor point accuracy. In decoding, an optimization-based multilateration method surpasses traditional approaches in landmark localization. Additionally, the *Equivalent Single-Step Parallel Computation* algorithm significantly improves computational efficiency, enabling real-time FLD applications. Extensive evaluations on five benchmark datasets demonstrate that POPoS consistently outperforms existing FLD methods, particularly in low-resolution scenarios.

Acknowledgments

Zhi-Qi Cheng's research for this project was supported in part by the US Department of Transportation, Office of the Assistant Secretary for Research and Technology, under the University Transportation Center Program (Federal Grant Number 69A3551747111). This work was also supported in part by the National Natural Science Foundation of China (Grant No. 62372387) and the Key R&D Program of Guangxi Zhuang Autonomous Region, China (Grant Nos. AB22080038 and AB22080039).

References

- Artizzu, X. P. B.; Perona, P.; and Dollár, P. 2013. Robust Face Landmark Estimation under Occlusion. In *ICCV*, 1513–1520.
- Bao, X.; Cheng, Z.-Q.; He, J.-Y.; Xiang, W.; Li, C.; Sun, J.; Liu, H.; Liu, W.; Luo, B.; Geng, Y.; et al. 2023. Key-PosS: Plug-and-Play Facial Landmark Detection through GPS-Inspired True-Range Multilateration. In *ACM Multimedia*, 5746–5755.
- Bulat, A.; Sanchez, E.; and Tzimiropoulos, G. 2021. Subpixel Heatmap Regression for Facial Landmark Localization. In *BMVC*, 422.
- Chandran, P.; Bradley, D.; Gross, M. H.; and Beeler, T. 2020. Attention-Driven Cropping for Very High Resolution Facial Landmark Detection. In *CVPR*.
- Cheng, Z.; Cheng, Z.-Q.; He, J.-Y.; Sun, J.; Wang, K.; Lin, Y.; Lian, Z.; Peng, X.; and Hauptmann, A. 2024a. Emotion-LLaMA: Multimodal Emotion Recognition and Reasoning with Instruction Tuning. In *arXiv preprint arXiv:2406.11161*.
- Cheng, Z.; Niu, F.; Lin, Y.; Cheng, Z.-Q.; Zhang, B.; and Peng, X. 2024b. MIPS at SemEval-2024 Task 3: Multimodal Emotion-Cause Pair Extraction in Conversations with Multimodal Language Models. In *SemEval*.
- Cheng, Z.; Tu, S.; Huang, D.; Li, M.; Peng, X.; Cheng, Z.-Q.; and Hauptmann, A. G. 2024c. SZTU-CMU at MER2024: Improving Emotion-LLaMA with Conv-Attention for Multimodal Emotion Recognition. In *MRAC*.
- Cheng, Z.-Q.; Dai, Q.; Li, H.; Song, J.; Wu, X.; and Hauptmann, A. G. 2022. Rethinking spatial invariance of convolutional networks for object counting. In *CVPR*, 19638–19648.
- Cheng, Z.-Q.; Li, J.-X.; Dai, Q.; Wu, X.; and Hauptmann, A. G. 2019a. Learning spatial awareness to improve crowd counting. In *ICCV*, 6152–6161.
- Cheng, Z.-Q.; Li, J.-X.; Dai, Q.; Wu, X.; He, J.-Y.; and Hauptmann, A. G. 2019b. Improving the Learning of Multicolumn Convolutional Neural Network for Crowd Counting. In *ACM Multimedia*, 1897–1906.
- Cheng, Z.-Q.; Liu, Y.; Wu, X.; and Hua, X.-S. 2016. Video ecommerce: Towards online video advertising. In *ACM Multimedia*, 1365–1374.
- Cheng, Z.-Q.; Wu, X.; Liu, Y.; and Hua, X.-S. 2017a. Video ecommerce++: Toward large scale online video advertising. *IEEE Trans. Multimedia*, 19(6): 1170–1183.
- Cheng, Z.-Q.; Wu, X.; Liu, Y.; and Hua, X.-S. 2017b. Video2shop: Exact matching clothes in videos to online shopping images. In *CVPR*, 4048–4056.
- Cheng, Z.-Q.; Zhang, H.; Wu, X.; and Ngo, C.-W. 2017c. On the selection of anchors and targets for video hyperlinking. In *ACM Multimedia Retrieval*, 287–293.
- Dapogny, A.; Cord, M.; and Bailly, K. 2019. DeCaFA: Deep Convolutional Cascade for Face Alignment in the Wild. In *ICCV*, 6892–6900.
- Fard, A. P.; and Mahoor, M. H. 2022. ACR loss: Adaptive coordinate-based regression loss for face alignment. In *ICCV*, 1807–1814. IEEE.
- Feng, Z.; Kittler, J.; Awais, M.; Huber, P.; and Wu, X. 2018. Wing Loss for Robust Facial Landmark Localisation With Convolutional Neural Networks. In *CVPR*, 2235–2245.
- He, J.-Y.; Cheng, Z.-Q.; Li, C.; Sun, J.; He, Q.; Xiang, W.; Chen, H.; Lan, J.-P.; Lin, X.; Zhu, K.; et al. 2024. MetaDesigner: Advancing Artistic Typography through AI-Driven, User-Centric, and Multilingual WordArt Synthesis. In *arXiv preprint arXiv:2406.19859*.
- He, J.-Y.; Cheng, Z.-Q.; Li, C.; Sun, J.; Xiang, W.; Lin, X.; Kang, X.; Jin, Z.; Hu, Y.; Luo, B.; et al. 2023. WordArt Designer: User-Driven Artistic Typography Synthesis using Large Language Models. In *EMNLP Industry Track*.
- Huang, S.; Li, X.; Cheng, Z.-Q.; Zhang, Z.; and Hauptmann, A. 2020a. Stacked pooling for boosting scale invariance of crowd counting. In *ICASSP*, 2578–2582.
- Huang, S.; Zhu, Z.; Guo, F.; and Huang, G. 2020b. The Devil Is in the Details: Delving Into Unbiased Data Processing for Human Pose Estimation. In *CVPR*, 5699–5708.
- Huang, Y.; Chen, X.; Kim, J.; Yang, H.; Li, C.; Yang, J.; and Chen, D. 2023. FreeEnricher: Enriching Face Landmarks without Additional Cost. In *AAAI*, 962–970.
- Jin, H.; Liao, S.; and Shao, L. 2021. Pixel-in-Pixel Net: Towards Efficient Facial Landmark Detection in the Wild. *Int. J. Comput. Vis.*, 129(12): 3174–3194.
- Jin, S.; Xu, L.; Xu, J.; Wang, C.; Liu, W.; Qian, C.; Ouyang, W.; and Luo, P. 2020. Whole-Body Human Pose Estimation in the Wild. In *ECCV*, volume 12354, 196–214.
- Khabaralakh, K.; and Koriashkina, L. 2021. Fast Facial Landmark Detection and Applications: A Survey. *arXiv preprint arXiv:2101.10808*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Köstinger, M.; Wohlhart, P.; Roth, P. M.; and Bischof, H. 2011. Annotated Facial Landmarks in the Wild: A large-scale, real-world database for facial landmark localization. In *ICCV Workshops*, 2144–2151.
- Kumar, A.; Marks, T. K.; Mou, W.; Wang, Y.; Jones, M.; Cherian, A.; Koike-Akino, T.; Liu, X.; and Feng, C. 2020. LUVLi Face Alignment: Estimating Landmarks' Location, Uncertainty, and Visibility Likelihood. In *CVPR*, 8233–8243.
- Li, H.; Guo, Z.; Rhee, S.; Han, S.; and Han, J. 2022a. Towards Accurate Facial Landmark Detection via Cascaded Transformers. In *CVPR*, 4166–4175.

- Li, J.; Jin, H.; Liao, S.; Shao, L.; and Heng, P.-A. 2022b. Repformer: Refinement pyramid transformer for robust facial landmark detection. In *IJCAI*.
- Li, M.; Li, H.; Cheng, Z.-Q.; Dong, Y.; Zhou, Y.; He, J.-Y.; Dai, Q.; Mitamura, T.; and Hauptmann, A. G. 2024. Human-Aware Vision-and-Language Navigation: Bridging Simulation to Reality with Dynamic Human Interactions. In *arXiv preprint arXiv:2406.19236*.
- Li, W.; Lu, Y.; Zheng, K.; Liao, H.; Lin, C.; Luo, J.; Cheng, C.; Xiao, J.; Lu, L.; Kuo, C.; and Miao, S. 2020. Structured Landmark Detection via Topology-Adapting Deep Graph Learning. In *ECCV*, volume 12354, 266–283.
- Lv, J.; Shao, X.; Xing, J.; Cheng, C.; and Zhou, X. 2017. A Deep Regression Architecture with Two-Stage Re-initialization for High Performance Facial Landmark Detection. In *CVPR*, 3691–3700.
- Newell, A.; Yang, K.; and Deng, J. 2016. Stacked Hourglass Networks for Human Pose Estimation. In *ECCV*, volume 9912, 483–499.
- Nibali, A.; He, Z.; Morgan, S.; and Prendergast, L. A. 2018. Numerical Coordinate Regression with Convolutional Neural Networks. *CoRR*, abs/1801.07372.
- Papandreou, G.; Zhu, T.; Kanazawa, N.; Toshev, A.; Tompson, J.; Bregler, C.; and Murphy, K. 2017. Towards Accurate Multi-person Pose Estimation in the Wild. In *CVPR*, 3711–3719.
- Qian, S.; Sun, K.; Wu, W.; Qian, C.; and Jia, J. 2019. Aggregation via Separation: Boosting Facial Landmark Detector With Semi-Supervised Style Translation. In *ICCV*, 10152–10162.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, volume 9351, 234–241.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.*, 115(3): 211–252.
- Sagonas, C.; Antonakos, E.; Tzimiropoulos, G.; Zafeiriou, S.; and Pantic, M. 2016. 300 Faces In-The-Wild Challenge: database and results. *Image Vis. Comput.*, 47: 3–18.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep High-Resolution Representation Learning for Human Pose Estimation. In *CVPR*, 5693–5703.
- Tai, Y.; Liang, Y.; Liu, X.; Duan, L.; Li, J.; Wang, C.; Huang, F.; and Chen, Y. 2019. Towards Highly Accurate and Stable Face Alignment for High-Resolution Videos. In *AAAI*, 8893–8900.
- Tang, Z.; Peng, X.; Geng, S.; Wu, L.; Zhang, S.; and Metaxas, D. N. 2018. Quantized Densely Connected U-Nets for Efficient Landmark Localization. In *ECCV*, 348–364.
- Tu, S.; Dai, Q.; Cheng, Z.-Q.; Hu, H.; Han, X.; Wu, Z.; and Jiang, Y.-G. 2024a. Motioneditor: Editing video motion via content-aware diffusion. In *CVPR*, 7882–7891.
- Tu, S.; Dai, Q.; Zhang, Z.; Xie, S.; Cheng, Z.-Q.; Luo, C.; Han, X.; Wu, Z.; and Jiang, Y.-G. 2024b. MotionFollower: Editing Video Motion via Lightweight Score-Guided Diffusion. In *arXiv preprint arXiv:2405.20325*.
- Wan, J.; Liu, J.; Zhou, J.; Lai, Z.; Shen, L.; Sun, H.; Xiong, P.; and Min, W. 2023. Precise Facial Landmark Detection by Reference Heatmap Transformer. *IEEE Trans. Image Process.*, 32: 1966–1977.
- Wang, X.; Bo, L.; and Li, F. 2019. Adaptive Wing Loss for Robust Face Alignment via Heatmap Regression. In *ICCV*, 6970–6980.
- Wu, W.; Qian, C.; Yang, S.; Wang, Q.; Cai, Y.; and Zhou, Q. 2018. Look at Boundary: A Boundary-Aware Face Alignment Algorithm. In *CVPR*, 2129–2138.
- Wu, Y.; and Ji, Q. 2018. Facial Landmark Detection: a Literature Survey. *arXiv preprint arXiv:1805.05563*.
- Xia, J.; Qu, W.; Huang, W.; Zhang, J.; Wang, X.; and Xu, M. 2022. Sparse local patch transformer for robust face alignment and landmarks inherent relation learning. In *CVPR*, 4052–4061.
- Xu, C.; Liu, Y.; Xing, J.; Wang, W.; Sun, M.; Dan, J.; Huang, T.; Li, S.; Cheng, Z.-Q.; Tai, Y.; et al. 2024a. FaceChain-ImagineID: Freely Crafting High-Fidelity Diverse Talking Faces from Disentangled Audio. In *CVPR*, 1292–1302.
- Xu, C.; Sun, M.; Cheng, Z.-Q.; Wang, F.; Liu, Y.; Sun, B.; Huang, R.; and Hauptmann, A. 2024b. Combo: Co-speech holistic 3D human motion generation and efficient customizable adaptation in harmony. *arXiv preprint arXiv:2408.09397*.
- Zeng, W.; et al. 2023. 3D-Aware Facial Landmark Detection via Multi-View Consistent Training on Synthetic Data. In *CVPR*.
- Zhang, F.; Zhu, X.; Dai, H.; Ye, M.; and Zhu, C. 2020. Distribution-Aware Coordinate Representation for Human Pose Estimation. In *CVPR*, 7091–7100.
- Zhang, J.; Cheng, Z.-Q.; Wu, X.; Li, W.; and Qiao, J.-J. 2022. Crossnet: Boosting crowd counting with localization. In *ACM Multimedia*, 6436–6444.
- Zhou, Z.; Li, H.; Liu, H.; Wang, N.; Yu, G.; and Ji, R. 2023. STAR Loss: Reducing Semantic Ambiguity in Facial Landmark Detection. In *CVPR*, 15475–15484.
- Zhu, M.; Shi, D.; Zheng, M.; and Sadiq, M. 2019. Robust Facial Landmark Detection via Occlusion-Adaptive Deep Networks. In *CVPR*, 3486–3496.
- Zou, X.; Zhong, S.; Yan, L.; Zhao, X.; Zhou, J.; and Wu, Y. 2019. Learning Robust Facial Landmark Detection via Hierarchical Structured Ensemble. In *ICCV*, 141–150.