



PAL-Net: A point-wise CNN with patch-attention for 3D anatomical facial landmark localization

Ali Shadman Yazdi ^a,* Annalisa Cappella ^{b,c}, Benedetta Baldini ^a, Riccardo Solazzo ^b, Gianluca Tartaglia ^{d,e}, Chiarella Sforza ^b, Giuseppe Baselli ^a

^a Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy

^b Department of Biomedical Sciences for Health, University of Milan, Milan, Italy

^c U.O. Laboratory of Applied Morphology, IRCCS Policlinico San Donato, 20097 San Donato Milanese, Italy

^d Department of Biomedical, Surgical and Dental Sciences, University of Milan, Milan, Italy

^e Fondazione IRCCS Cà Granda, Ospedale Maggiore Policlinico, 20122 Milan, Italy

ARTICLE INFO

Keywords:

3D facial landmark localization

Anatomical landmarks

Patch-based CNN

Point-wise convolution

Facial morphometrics

ABSTRACT

Manual annotation of anatomical landmarks on 3D facial scans is a time-consuming and expertise-dependent task, yet it remains critical for clinical assessments, morphometric analysis, and craniofacial research. While several deep learning methods have been proposed for facial landmark localization, most focus on pseudo-landmarks or require complex input representations, limiting their clinical applicability. This study presents a fully automated deep learning pipeline (PAL-Net) for localizing 50 anatomical landmarks on facial models acquired via stereo-photogrammetry. The method combines coarse alignment, region-of-interest filtering, and an initial landmark approximation with a patch-based pointwise CNN enhanced by attention mechanisms. Trained and evaluated on 214 annotated scans from healthy adults, PAL-Net achieved a mean localization error of 3.686 mm and preserved relevant anatomical distances with an average error of 2.822 mm. While the geometric error exceeds expert intra-observer variability, the distance-wise error maintains structural integrity sufficient for high-throughput anthropometric analysis. To assess generalization, the model was further evaluated on 700 subjects from the FaceScape dataset, achieving a mean localization error of 0.41 mm and a distance error of 0.38 mm. Comparing with existing methods, PAL-Net offers a favorable trade-off between accuracy and computational cost. While performance degrades in regions with poor mesh quality (e.g., ears, hairline), the method demonstrates consistent accuracy across most anatomical regions. PAL-Net generalizes effectively across datasets and facial regions, outperforming existing methods in both point-wise and structural evaluations. It provides a lightweight, scalable solution for high-throughput 3D anthropometric analysis, with potential to support clinical workflows and reduce reliance on manual annotation. Source code can be accessed at <https://github.com/AliShadman/PAL-Net-A-Point-Wise-CNN-with-Patch-Attention>.

1. Introduction

Precise localization of anatomical facial landmarks is fundamental to clinical craniofacial analysis, guiding surgical planning, outcome assessment, and dysmorphology diagnosis [1–4]. Beyond pathology, 3D anthropometry is essential for studying growth trajectories and population variations [5]. While modern stereo-photogrammetry systems deliver high-resolution 3D meshes comparable to those from laser scanning [6], downstream analysis still rely heavily on accurate landmarking. However, facial landmarks fall into two distinct categories: anatomical and pseudo-landmarks. While pseudo-landmarks are algorithmically defined by geometric or statistical properties of the

facial images and widely used in general computer vision, anatomical landmarks are standard for medical anthropometry. They correspond to specific biological structures [7] and have consistent definitions across different individuals [8].

Despite the advancements in digitized anthropometry, identifying these anatomical points remains a bottleneck. Manual annotation is labor-intensive and prone to inter-observer variability, particularly in regions with few distinct geometric features [9,10]. Existing automated methods often focus on pseudo-landmarks or require computationally expensive inputs, limiting their direct clinical applicability. To address this, we introduce PAL-Net, a deep learning framework tailored for the

* Correspondence to: Via Camillo Golgi, 39, 20133 Milano MI, Italy.

E-mail address: ali.shadman@polimi.it (A.S. Yazdi).

<https://doi.org/10.1016/j.imu.2025.101729>

Received 3 October 2025; Received in revised form 15 December 2025; Accepted 16 December 2025

Available online 17 December 2025

2352-9148/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

efficient and accurate localization of 50 anatomical landmarks on 3D facial scans. The main contributions of this study are summarized as follows:

- We propose PAL-Net, a lightweight patch-based point-wise CNN that integrates a global attention mechanism to capture both local geometry and global facial structure.
- We introduce a composite loss function that explicitly penalizes deviations in inter-landmark distances. This structure-preserving design ensures that the predicted geometry maintains the anthropometric consistency required for clinical measurements.
- We design the framework to operate directly on raw, unstructured point clouds, eliminating the need for computationally intensive preprocessing.
- We target a comprehensive set of 50 clinically validated anatomical landmarks overcoming the limitations of previous methods restricted to pseudo-landmarks.
- We demonstrate that PAL-Net achieves a favorable trade-off, providing accuracy that is sufficient for anthropometric use while significantly reducing computational cost compared to state-of-the-art methods.
- We provide a robust evaluation on both a clinically annotated dataset (LAFAS) and a large-scale public dataset (FaceScape), validating the model's generalization capabilities.

The remainder of this paper is organized as follows. Section 2 outlines the problem formulation and details the proposed methodology, including data preprocessing and the PAL-Net architecture. Section 3 describes the datasets, training protocols, and experimental setup. Section 4 presents the comparative and anatomical performance analysis, including ablation studies. Finally, Section 5 discusses limitations and concludes the study with directions for future work.

1.1. Related works

Approaches for 3D anatomical facial landmark localization have progressed from classical geometric algorithms to modern learning-based techniques. Initially, traditional methods, predating machine learning, focused on hand-crafted geometric feature extraction and model fitting. These techniques identify landmarks using local surface properties such as principal curvatures, shape indices, or geodesic distances [11]. Additionally, some methods employ statistical shape models, like 3D Active Shape Models (ASMs), or use deformable models to fit a template to a target [12]. In parallel, dense correspondence methods attempt to map the entire facial surface by modeling coherent local deformations; for example, boosting local shape matching [13] and landmark-free registration [14]. While these dense correspondence techniques can implicitly localize landmarks via full-template alignment, they address a broader and more computationally intensive registration problem. In contrast, our method focuses on directly and efficiently localizing a sparse set of clinically relevant anatomical points, optimizing for high-throughput anthropometric analysis.

To address the need for efficiency and robustness, recent studies have investigated the application of machine learning (ML) and deep learning (DL) techniques specifically for landmark localization. Since pseudo-landmarks are algorithmically determined, they exhibit high consistency with minimal variability; therefore, most studies on facial landmark localization develop their models using these landmarks, as their consistency often leads to improved results. Creusot et al. [15] used machine learning to annotate 14 facial landmarks on 3D models by manually extracting features such as vector normals, neighboring point data, and principal curvatures. Their approach, combining offline training and online detection, achieved prediction errors ranging from 2.5 to 10 mm on non-anatomical landmarks. O'Sullivan [16] demonstrated the feasibility of adapting 2D methods to 3D facial landmark localization by extending Convolutional Pose Machines (CPMs) [17]

for 3D facial landmark localization, by refining heat-maps using PointNet++ [18] architecture. Their approach demonstrates the feasibility of extending 2D methods to 3D facial landmark localization. Paulsen et al. [19] achieved a localization error of 2.42 mm using a multi-view approach for facial landmark identification on the BU-3DFE datasets [20], comprising 83 non-anatomical facial landmarks. Although the results are very accurate, the methodology requires significantly high computational power. It involves rendering the 3D facial model from multiple views, generating multiple 2D images of the face, and applying 2D methods for facial landmark localization. Wang et al. [21] used Graph Convolutional Networks (GCN) with PAConv [22] for 3D facial landmark localization, learning self-attention, and refining affine transformations. Their method achieved a 2 mm localization error, outperforming others on the same dataset. Burger et al. [23] employed a two-stage stratified graph convolutional network model for facial landmark detection on 3D data(2S-SGCN). The first stage uses a stratified graph convolutional network (SGCN) to detect landmark regions by combining global and local graph representations. The second stage refines these predictions with the MSE-over-mesh method to accurately locate landmarks directly on the mesh. This approach, while being computationally efficient, achieves an average localization accuracy of approximately 0.371 mm on 3D scans from the FaceScape dataset [24].

The mentioned studies have applied their methodology to pseudo-landmarks due to their consistency. However, fewer studies have applied various methods to localize anatomical landmarks, due to the limited availability of annotated training data, especially in medical imaging contexts [25]. This scarcity makes it difficult to train models effectively, leading to lower accuracy in landmark detection. Chong et al. [26] developed a U-NET-based deep learning algorithm [27] for automated anatomical landmark detection on 3D facial images. The method involves stacking two U-NETs for coarse and fine feature extraction, followed by back-projection to accurately localize 20 landmarks. The model was validated on healthy subjects, acromegaly patients, and localized scleroderma patients, achieving an average normalized mean error (NME) of 1.4 mm for healthy cases, 2.2 mm for scleroderma patients, and 2.8 mm for acromegaly patients. Guo et al. [28] developed a method for automatic 3D facial landmark annotation using a combination of anatomical landmarks and pseudo-landmarks for automatic facial image registration. They began with the manual annotation of six key anatomical landmarks, which were then automatically localized using PCA-based feature recognition. In addition to these anatomical landmarks, 11 pseudo-landmarks were heuristically identified based on geometric relations and texture constraints. The study reported that the localization error for automatic landmark annotation ranged between 1.0 mm and 3.6 mm with factors like facial hair affecting the accuracy in some regions. Berends et al. [29] introduced a fully automated pipeline for localizing ten soft-tissue landmarks on 3D facial meshes using a two-stage DiffusionNet architecture. Their method combines spectral surface features (HKS), pose normalization, and non-rigid registration to achieve high accuracy (1.69 ± 1.15 mm) on a dataset of 2897 subjects. More recently, sophisticated deep learning architectures such as Transformer-based architectures have been explored for 3D position and anatomical landmark localization. For instance, Zhang et al. [30] proposed a cascade Point Transformer for landmark localization on 3D human scans, while Kubík et al. [31] applied a Point Transformer module for anatomical landmark detection on 3D dental images. These approaches, while powerful, often introduce significant computational complexity, motivating the need for more lightweight solutions. Addressing prior studies' limitations such as manual feature extraction, limited anatomical scope, and high computational requirements, this work proposes a lightweight deep learning-based approach to fully automate 3D anatomical facial landmark annotation. While existing methods often focus on a small subset of identifiable anatomical landmarks [32], typically focusing on fewer than 20, our approach targets 50 clinically validated anatomical landmarks across diverse facial regions, including areas known to be more

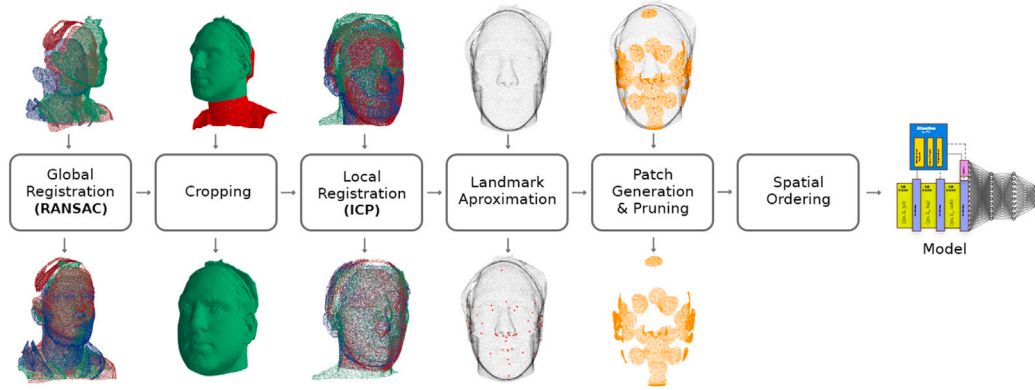


Fig. 1. Preprocessing pipeline for 3D facial data used in the study. The pipeline consists of coarse registration, cropping, local registration.

challenging to localize. The proposed model was trained and evaluated on a dataset of 214 stereo-photogrammetric facial models with expert-annotated anatomical landmarks. The proposed system uses minimal preprocessing and operates directly on point clouds without expensive mesh fitting or spectral descriptors. The methodology is scalable for real-world clinical and research settings, aiming to reduce manual annotation time, improve reproducibility, and provide a standardized high-resolution tool for both clinical and research applications in 3D anthropometric analysis.

2. Methods

This section describes the proposed pipeline for fully automated anatomical landmark localization on 3D facial meshes acquired via stereo-photogrammetry. The approach is developed and evaluated using the LAFAS dataset, a curated collection of 3D facial scans annotated with 50 clinically validated anatomical landmarks (see Section 3.2 for details). The main stages of the pipeline include: preprocessing and rigid alignment, coarse landmark approximation via population-based templates, prediction using a patch-based point-wise CNN with attention (PAL-Net), and evaluation based on both point-wise and inter-landmark distance errors. An overview of the preprocessing steps is shown in Fig. 1.

2.1. Preliminaries

Let a 3D facial model be represented as a point cloud $\mathcal{F} = \{\mathbf{f}_i\}_{i=1}^{n_f} \subset \mathbb{R}^3$, where n_f is the total number of vertices in the scan and each vertex \mathbf{f}_i corresponds to its Cartesian coordinates (x, y, z) . The objective is to predict the 3D coordinates of n_l predefined anatomical landmarks, for which the ground truth locations are denoted as the set $\mathcal{L} = \{\mathbf{l}_k\}_{k=1}^{n_l} \subset \mathbb{R}^3$, where each \mathbf{l}_k is a point on the facial surface, $\mathcal{L} \subset \mathcal{F}$. Our model, PAL-Net, learns a mapping function M to produce a set of predicted landmarks, $\hat{\mathcal{L}} = M(\mathcal{F})$, that minimizes a loss function $\mathcal{J}(\mathcal{L}, \hat{\mathcal{L}})$. To achieve this, our pipeline first generates an initial coarse approximation of the landmarks, $\tilde{\mathcal{L}}_{\text{fit}} \in \mathbb{R}^{n_l \times 3}$, via atlas-based alignment. From these approximations, we extract n_l local patches, where each patch $\mathcal{P}_k \in \mathbb{R}^{K \times 3}$ consists of the K nearest points to its corresponding approximated landmark. The final input to our network is the collection of all patches, $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_{n_l}\}$. The model then learns to refine these patches to predict the final, accurate landmark locations $\hat{\mathcal{L}}$.

2.2. Data preprocessing

First, to ensure computational efficiency without sacrificing anatomical accuracy, each 3D facial model was resampled using a rejection-sampling strategy. This resulted in a low-resolution representation containing approximately 10,000 evenly distributed surface points,

denoted as $\mathcal{F}_{\text{src}}^{\text{low}}$. Consequently, our pipeline utilizes two representations for each subject: the derived $\mathcal{F}_{\text{src}}^{\text{low}}$ for robust global registration, and the original high-resolution mesh ($\mathcal{F}_{\text{src}}^{\text{og}}$, $\approx 100,000$ vertices) for detailed local alignment. To ensure consistent anatomical alignment, we implemented a multi-stage registration strategy using these two resolutions. We began with a global registration step using Fast Point Feature Histograms (FPFH) and RANSAC to align the source model $\mathcal{F}_{\text{src}}^{\text{low}}$ against a fixed reference template, denoted as \mathcal{M}^{low} . This step computes a coarse rigid alignment, producing a transformation matrix $\mathbf{T}_{\text{coarse}}$, representing rigid body transformations within the Special Euclidean group $SE(3)$.

$$\mathbf{T}_{\text{coarse}} \in SE(3), \quad \text{such that} \quad \mathcal{F}_{\text{aligned}}^{\text{low}} = \mathbf{T}_{\text{coarse}} \cdot \mathcal{F}_{\text{src}}^{\text{low}}$$

This transformation is then applied to the original full-resolution facial mesh to bring it into the same coarse-aligned coordinate frame:

$$\mathcal{F}_{\text{aligned}}^{\text{og}} = \mathbf{T}_{\text{coarse}} \cdot \mathcal{F}_{\text{src}}^{\text{og}}$$

The result provided a coarse alignment, which are then applied to the high-resolution models, yielding an anatomically consistent set of high-resolution facial models. To enhance subsequent local registration, irrelevant background geometry was removed using a Region-of-Interest (ROI) filter. Defined as a bounding box $\mathcal{B}_{\text{ref}} \subset \mathbb{R}^3$ on the reference, the filter cropped the mesh by discarding vertices $\mathbf{f}_i \notin \mathcal{B}_{\text{ref}}$, yielding the subset:

$$\mathcal{F}_{\text{ROI}} = \left\{ \mathbf{f}_i \in \mathcal{F}_{\text{aligned}}^{\text{og}} \mid \mathbf{f}_i \in \mathcal{B}_{\text{ref}} \right\}$$

Cropping was performed at this stage solely to optimize the registration process, as the presence of irrelevant regions would have negatively affected both its accuracy and efficiency. For fine alignment, a rigid transformation $\mathbf{T}_{\text{ICP}} \in SE(3)$ was estimated via Iterative Closest Point (ICP) between \mathcal{F}_{ROI} and the reference, ensuring submillimeter correspondence. The final transformation composes the coarse and fine alignments, ensuring that no anatomical information was lost due to the temporary cropping.

$$\mathbf{T}_{\text{final}} = \mathbf{T}_{\text{ICP}} \cdot \mathbf{T}_{\text{coarse}}$$

Following registration, the dataset was partitioned into a training set (80%) and a validation set (20%) to balance data availability and evaluation needs. Given the limited sample size, no separate test set was created.

To establish a population-based atlas, we computed the mean landmarks $\tilde{\mathcal{L}}_{\text{mean}}$ by averaging the training set of N samples: $\tilde{\mathcal{L}}_{\text{mean}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}^{(i)}$. This template was then transferred to each individual mesh $\mathcal{F}_{\text{aligned}}^{\text{og}}$ to provide an initial landmark estimate. Specifically, each template landmark $\tilde{\mathbf{l}}_k \in \tilde{\mathcal{L}}_{\text{mean}}$ was projected onto the mesh surface by selecting the nearest vertex in Euclidean space i.e., $\tilde{\mathbf{l}}_k^{\text{fit}} = \arg \min_{\mathbf{f}_i \in \mathcal{F}_{\text{aligned}}^{\text{og}}} \|\tilde{\mathbf{l}}_k - \mathbf{f}_i\|_2$. The resulting set of fitted estimates, $\tilde{\mathcal{L}}_{\text{fit}} \in \mathbb{R}^{n_l \times 3}$, serves as an initial approximation of the landmark positions.

These coordinates center the local patches used as input for the prediction model, which then refines these initial projections to determine the precise anatomical locations. For each landmark $\hat{\mathbf{l}}_k^{\text{fit}} \in \mathbb{R}^{1 \times 3}$, a localized patch is extracted from the full facial point cloud $\mathcal{F}^{\text{og}}_{\text{aligned}}$. Two strategies can be employed to generate these patches. In the first approach, the K nearest points in the aligned facial model $\mathcal{F}^{\text{og}}_{\text{aligned}}$ are selected for each fitted landmark $\hat{\mathbf{l}}_k^{\text{fit}}$, forming a local patch $\mathcal{P}_k \in \mathbb{R}^{K \times 3}$ centered around that landmark. In the second, all points within a fixed radius D from the fitted landmark $\hat{\mathbf{l}}_k^{\text{fit}}$ are selected and then uniformly resampled to K points to ensure a consistent patch size. In both cases, the resulting collection of patches for all landmarks is denoted as $\mathcal{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_{n_l}\} \in \mathbb{R}^{n_l \times K \times 3}$.

To ensure consistency, the K points within each patch are further sorted by their Euclidean distance to the origin of the global reference coordinate system $\mathbf{o} \in \mathbb{R}^{1 \times 3}$, defined near the nasal region of the face. This structured and repeatable input format improves the model's ability to capture spatial relationships while reducing training instability that may arise from arbitrary point orderings. The choice of K or D directly affects patch granularity and is further evaluated in the ablation experiments presented in the Results section.

2.3. Patch-Attention Landmark Network (PAL-Net)

The proposed Patch-Attention Landmark Network (PAL-Net), demonstrated in Fig. 2, refines 3D facial landmark positions by leveraging both local and global spatial information extracted from structured point cloud patches. The model operates on a structured tensor $\mathcal{X} \in \mathbb{R}^{m \times n \times K \times 3}$, where m is the number of subjects (or batch size), n is the number of landmarks per subject, K is the number of points in each patch, and 3 corresponds to the Cartesian coordinates (x, y, z) . To extract features from this spatially ordered representation, PAL-Net employs a series of convolutional layers using point-wise 1×1 convolutions [33]. Each convolutional block consists of two successive point-wise operations applied independently to each point within a patch:

$$\mathcal{X}'_{ijkf} = \sigma \left(\sum_{c=1}^3 W_{cf}^{(l)} \cdot \mathcal{X}_{ijkc} + b_f^{(l)} \right)$$

where $W^{(l)} \in \mathbb{R}^{3 \times f}$ and $b^{(l)} \in \mathbb{R}^f$ denote the learnable parameters of the l th convolutional layer, f is the number of output feature channels, and σ is the activation function, which in our case is a Rectified Linear Unit (ReLU). The model comprises three sequential convolutional blocks, each consisting of two point-wise 1×1 convolutional layers with ReLU activations. The first block employs 32 filters, followed by a max pooling operation along the patch dimension (K), reducing the number of points by a factor of 5. The second block increases the feature dimensionality to 64 and applies an identical pooling strategy. The third block uses 128 filters and a max pooling operation by a factor of 4. Pooling is selectively applied only along the third axis (K), while the second axis (the number of landmarks, n) is left unchanged to preserve the structural organization of landmark-wise patches. This strategy ensures that spatial resolution within each patch is progressively reduced, while maintaining the inter-landmark relationships established during the preprocessing stage.

Attention integration

To incorporate global context and enhance prediction accuracy, PAL-Net integrates an attention mechanism after each convolutional block. Prior to pooling, the feature maps $\mathcal{X}' \in \mathbb{R}^{m \times n \times K \times f}$ are reshaped into a 2D sequence $S \in \mathbb{R}^{m \times (n \cdot K) \times f}$, concatenating all points across all landmarks for each subject. An attention weight matrix $\mathcal{A} \in \mathbb{R}^{m \times (n \cdot K) \times 1}$ is computed as:

$$\mathcal{A} = \text{Softmax}(\tanh(SW + b))$$

where $W \in \mathbb{R}^{f \times 1}$ and $b \in \mathbb{R}^1$ are trainable parameters. The softmax operation ensures that the attention weights across all $n \cdot K$ points sum

to 1, producing a probabilistic weighting over the input sequence. The weighted global feature vector is then computed as $G_i = \sum_{j=1}^{n \cdot K} \mathcal{A}_{ij} \cdot S_{ij}$, resulting in $G \in \mathbb{R}^{m \times f}$. This global descriptor captures the most informative spatial features across the full facial geometry and is concatenated with the local features from the final convolutional block, forming a hybrid feature representation $H \in \mathbb{R}^{m \times n \times h}$, where h is the combined feature dimensionality. The final feature tensor H is processed by a multilayer perceptron (MLP) composed of three fully connected layers. The first layer maps each feature vector to 1024 dimensions, followed by a ReLU activation and dropout ($p=0.1$) regularization to introduce nonlinearity and prevent overfitting. The second fully connected layer also has 1024 units and uses a linear activation, allowing the model to adjust the features to the output scale. Finally, a third linear layer projects each feature vector to \mathbb{R}^3 , corresponding to the (x, y, z) coordinates of each anatomical landmark. The output is then reshaped to form the final predicted landmark positions, $\hat{\mathcal{L}} \in \mathbb{R}^{m \times n \times 3}$.

2.4. Post-processing

In practice, the predicted landmark coordinates $\hat{\mathcal{L}} \in \mathbb{R}^{n_l \times 3}$ may not precisely lie on the surface of the 3D facial mesh $\mathcal{F} \in \mathbb{R}^{n_f \times 3}$ due to prediction continuity, mesh sparsity, or corrupted surface regions. To address this, we apply a post-processing step to project each predicted point back onto the mesh. Two strategies are considered: the first projects each predicted point $\hat{\mathbf{l}}_k$ to its nearest surface vertex, i.e., $\hat{\mathbf{l}}_k^{\text{proj}} = \arg \min_{\mathbf{f}_i \in \mathcal{F}} \|\hat{\mathbf{l}}_k - \mathbf{f}_i\|_2$; the second computes the centroid of the K nearest vertices, i.e., $\hat{\mathbf{l}}_k^{\text{cent}} = \frac{1}{K} \sum_{\mathbf{f}_j \in \mathcal{N}_K(\hat{\mathbf{l}}_k)} \mathbf{f}_j$, which provides a more stable estimate in degraded regions of the mesh (e.g., around the ears or hairline) where the surface is incomplete or noisy. In such cases, the ground truth landmark itself may not lie directly on a valid surface vertex. Depending on dataset characteristics and mesh quality, either strategy may be preferable to ensure robust and anatomically plausible outputs.

3. Experimental analysis

3.1. Training procedure

The model was trained using a composite loss function designed to balance point-wise accuracy with anatomical consistency. This objective minimizes a localization loss, defined as the mean Euclidean distance between predicted and ground-truth landmarks, combined with a structural loss based on the mean absolute difference of their inter-landmark distances. The final loss was formulated as a weighted sum of the two components, with weights set to $\alpha = 0.6$ and $\beta = 0.4$, prioritizing precise localization while encouraging the preservation of relative distances between landmarks:

$$\mathcal{L}_{\text{total}} = \alpha \cdot \frac{1}{n} \sum_{k=1}^n \|\hat{\mathbf{l}}_k - \mathbf{l}_k\|_2 + \beta \cdot \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left| \|\hat{\mathbf{l}}_i - \hat{\mathbf{l}}_j\|_2 - \|\mathbf{l}_i - \mathbf{l}_j\|_2 \right|$$

$\hat{\mathbf{l}}_k$ and \mathbf{l}_k denote the predicted and ground truth coordinates of the k th landmark, respectively. Model weights were initialized using the Glorot Normal (Xavier normal) method to facilitate efficient gradient propagation during backpropagation. A fixed random seed was used to ensure reproducibility across multiple runs. Training was performed using the Adam optimizer with an initial learning rate of 1×10^{-3} and a batch size of 16. To improve convergence and mitigate overfitting, a ReduceLROnPlateau scheduler was employed. The scheduler monitored the validation loss at each epoch and, if no improvement was observed for 8 consecutive epochs, it reduced the learning rate by a factor of 0.5. Additionally, an early stopping mechanism was applied with a patience of 30 epochs, meaning that training would terminate if no further improvement in validation loss was observed for that number of epochs. The model weights corresponding to the epoch with the lowest observed validation loss (typically occurring around epoch 250) were retained and later reloaded to ensure that subsequent evaluations

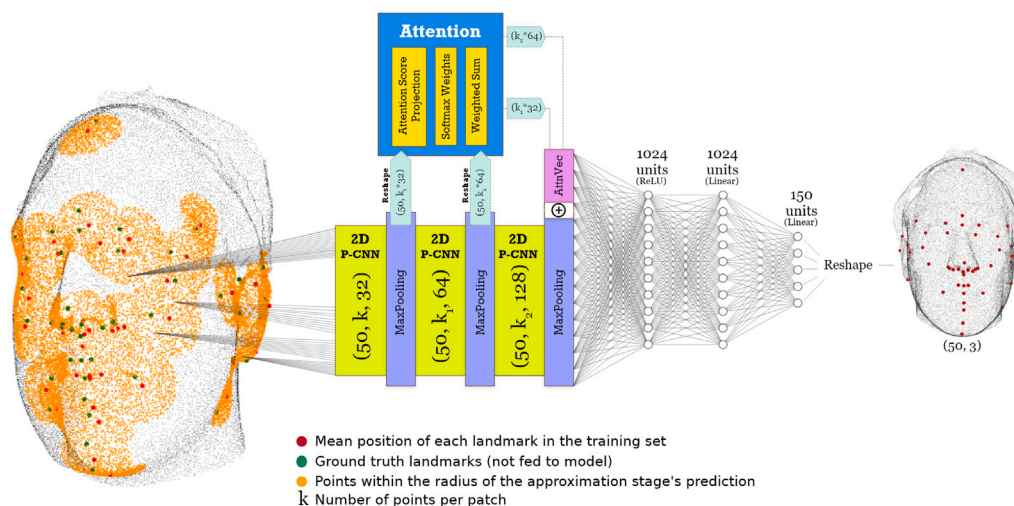


Fig. 2. Architecture overview of PAL-Net for predicting anatomical landmarks on the LAFAS dataset (50 landmarks). On the left, localized facial patches (orange) are extracted around approximated landmark positions (red) based on population-averaged coordinates, with ground truth annotations (green) shown for reference but not used as input. Patches are processed through a series of 2D point-wise CNN blocks with increasing feature depth and max pooling. Attention modules capture global context across all patches. The combined features are passed through fully connected layers to predict the 3D coordinates of 50 landmarks.

used the optimal parameters. Data preprocessing, patch generation, and landmark approximation were completed prior to training and remained unchanged throughout the training loop.

3.2. Datasets

3.2.1. LAFAS dataset

The LAFAS dataset served as the primary dataset for this study, and the development of the proposed model was specifically motivated by the need to automate anatomical landmark prediction on this dataset. More specifically, as it contains manually annotated anatomical landmarks with established relevance in real-world biomedical applications, its clinical relevance and standardized acquisition protocol provided an ideal foundation for training and evaluating a robust and generalizable anatomical landmark localization framework. The data were sourced from the LAFAS (Laboratory of Functional Anatomy of the Stomatognathic System of the Dipartimento di Scienze Biomediche per la Salute, Università degli Studi di Milano), which contains 214 3D facial scans of healthy subjects aged 18 to 49 years, both sexes. Excluded from this study were subjects with a previous history of craniofacial traumas, congenital anomalies, or craniofacial surgery. The 3D facial models were acquired using VECTRA M3, a fixed stereophotogrammetric device (Canfield Scientific Inc., Fairfield, NJ, USA), and VECTRA H1, a portable one. These instruments are employed routinely in the LAFAS and have demonstrated comparability and equivalence. In the study by [34], the validation of the Vectra 3D imaging system was verified. Each scan includes 50 facial landmarks manually annotated using a standardized protocol developed by [35]. These landmarks, shown in Fig. 3(a) and listed in Table 3, were manually annotated using a point-and-click interface by the skilled personnel of the LAFAS lab. To improve precision in manual landmark localization, specific landmarks were labeled with eyeliner on the face before acquisition. The 50 anatomical landmarks that cover various facial regions are employed in the LAFAS laboratory, in a variety of biomedical applications, including the assessment of facial asymmetry, morphometric facial characterization, and anthropometric analysis [36,37]. In these analyses, the relative distances between landmarks are typically utilized, rather than their absolute spatial positions. Many clinical and morphological assessments rely on geometric relationships and proportions between landmarks rather than their exact coordinate positions. To align the training objective with these downstream biomedical tasks,

the proposed model incorporates a loss component that explicitly penalizes deviations in inter-landmark distances. This design encourages the model not only to accurately localize points but also to preserve the geometric structure of the face, thereby enhancing the clinical relevance and interpretability of the predictions. To assess intra-observer reliability of the LAFAS dataset during the data annotation process, 20 random cases were selected, and the same operator annotated the same set of facial data twice within a two-week interval. This evaluation aimed to determine the consistency of the operator's annotations over time. The results of this reliability analysis are presented in Table 3, showcasing the mean difference in distance (mean Euclidean distance) between the two rounds of annotations.

3.2.2. FaceScope dataset

The second dataset used in this study was the FaceScope dataset, a publicly available 3D facial dataset introduced by Zhu et al. [38], Yang et al. [24]. It contains high-resolution 3D facial scans of over 960 subjects, collected using a multi-camera photogrammetry system. The dataset includes a wide range of facial expressions, head poses, and identity variations, making it a valuable resource for tasks such as 3D face modeling, expression synthesis, and geometry-aware facial analysis. For the purposes of this study, only neutral expression scans were used to ensure consistency with the training data and to provide a clearer evaluation of the model's generalization to new subjects under similar conditions. Out of the 847 available neutral scans, a subset of 700 subjects was used for training and evaluation. The 3D scans are accompanied by dense mesh correspondence as well as 68 annotated facial landmarks, shown in Fig. 3(b), primarily derived from automatic algorithms rather than manual annotation. While the dataset offers broad subject diversity and expression coverage, the landmark positions are not manually curated and do not follow a clinically validated annotation protocol, unlike the LAFAS dataset, which is tightly linked to biomedical applications and clinical use cases. The FaceScope dataset is primarily designed for general-purpose computer vision, and the annotated landmarks are not routinely used in clinical contexts and may not correspond precisely to anatomical reference points. For this reason, the FaceScope dataset was used solely for exploratory validation and generalization testing, while the model was developed and optimized based on the LAFAS dataset, which better reflects the clinical standards and requirements relevant to the target applications.

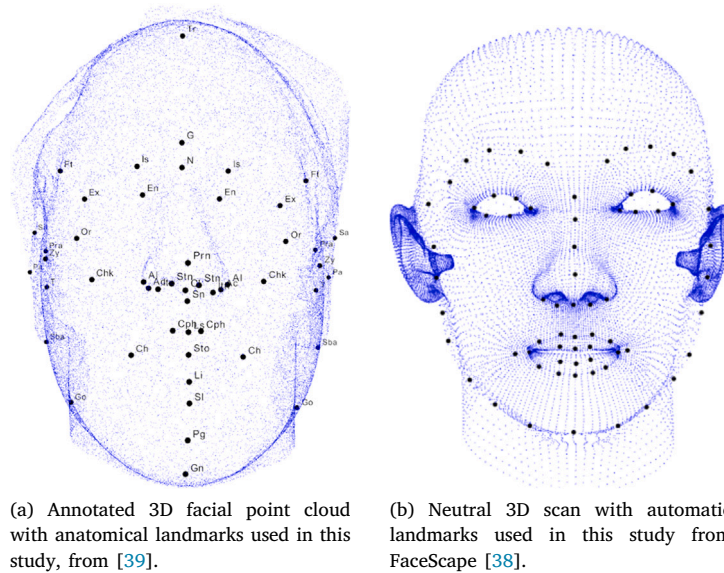


Fig. 3. Example 3D facial meshes with annotated landmarks used in this study.

4. Results and discussion

4.1. Comparative analysis with state-of-the-art methods

A comparison with existing methods across the two datasets is reported in Table 1. All models were trained and evaluated on the same hardware configuration using a Tesla T4 GPU and a batch size of 16 to ensure a fair and consistent benchmarking environment. For the Deep-MVLM model, results correspond to a single training fold, with the best-performing epoch selected from 100 total training epochs. For the LAFAS dataset, results for both 2S-SGCN and the proposed PAL-Net are reported as the mean prediction errors averaged over 5-fold cross-validation. In contrast, for the FaceScape dataset, evaluation was performed on a single fold due to its larger data volume. The increased sample size in FaceScape made full cross-validation computationally intensive and less critical for generalization testing. The hyperparameters of the 2S-SGCN model that was trained on FaceScape were the default configuration reported in the original study, specifically with a model depth of 20 and $k_{\text{knn}} = 16$. However, for the LAFAS dataset, this configuration did not yield optimal results; instead, a shallower architecture (depth = 4) and a larger local neighborhood ($k_{\text{knn}} = 32$) provided improved performance, highlighting the importance of dataset-specific tuning. In both datasets, the 2S-SGCN was trained for 500 epochs. These adjustments ensured that each model operated under conditions most suitable to its architecture and the dataset's specific characteristics, enabling a more meaningful performance comparison. All reported results for PAL-Net were obtained using localized facial patches constructed from the 1000 nearest vertices to each approximated landmark in the original facial meshes, trained on around 250 epochs with early stopping enabled. For the LAFAS dataset, the final prediction for each landmark was computed as the centroid of the $K = 10$ closest points in the predicted coordinate space, which helped mitigate local surface noise. In the FaceScape dataset, which contains denser, more complete facial meshes, the final landmark prediction was determined by selecting the surface point closest to the model's predicted coordinate. In addition to achieving lower mean errors, PAL-Net also demonstrated reduced standard deviation across both point-wise and distance-wise metrics, indicating more consistent predictions and improved robustness. This indicates more consistent predictions across subjects and landmarks, suggesting improved model stability and generalization. Another advantage of PAL-Net lies in

Table 1

Comparison of point-wise and distance-wise localization errors (in mm) for different methods on LAFAS and FaceScape datasets on the evaluation set corresponding to 20% of the total number of cases in each dataset. PAL-Net achieved the best accuracy with a significantly lower GPU memory requirement (2.5 GiB vs. 9–11 GiB).

Model	Metric	LAFAS (214 Subjects) (50 Landmarks)	FaceScape (700 Subjects) (68 Landmarks)
MVLM [19]	Point-wise (mm)	4.691 ± 3.486	5.254 ± 3.561
	Distance-wise (mm)	3.870 ± 3.396	3.975 ± 3.399
	Train Epoch Time	40 min	70 min
2S-SGCN [23]	Point-wise (mm)	3.689 ± 3.178	0.709 ± 0.944
	Distance-wise (mm)	2.973 ± 3.069	0.613 ± 0.778
	Train Epoch Time	41.4 s	81 s
PAL-Net (Ours)	Point-wise (mm)	3.686 ± 2.306	0.410 ± 0.634
	Distance-wise (mm)	2.822 ± 2.326	0.380 ± 0.536
	Train Epoch Time	3.32 s	7.51 s

its lightweight training characteristics, despite achieving competitive accuracy across both datasets. PAL-Net requires substantially less computational time per training epoch than alternative approaches, and this efficiency is primarily due to the model's compact architecture and the use of pointwise convolutional operations, which are well-suited for localized patch-based learning. Fig. A.6 shows the per-landmark localization error for each landmark of the LAFAS dataset, while Fig. A.7 displays the Pal-Net landmark predictions compared against the ground-truth values.

To evaluate the computational efficiency of PAL-Net in comparison with existing approaches, we measured the average runtime per subject on the LAFAS dataset, broken down into three stages: Input Preparation, Model Inference, and Final Prediction, as reported in Table 2. These measurements allow a detailed comparison of the end-to-end runtime characteristics of each method under realistic conditions. All evaluations were performed using an NVIDIA Tesla T4 GPU, and I/O-related delays (e.g., file imports) were explicitly excluded to focus on model-relevant computations. Inference timing was performed with a batch size of 1, and all other times are averaged per subject to ensure subject-level precision. For PAL-Net, input preparation includes generating localized point cloud patches and sorting the points; inference covers the model's forward pass; and the final prediction step projects

Table 2

Average runtime per subject (in seconds) for each model, broken down into input preparation, inference, and post-processing (final prediction). Manual annotation time by an expert is included for comparison.

Model	Input preparation	Model inference	Final prediction	Total time
Manual Annotation	–	–	–	6–7 min
MVLM [19]	0.431 s	11.552 s	5.101 s	12.084 s
2S-SGCN [23]	0.062 s	0.102 s	0.042 s	0.206 s
PAL-Net (Ours)	0.207 s	0.005 s	0.092 s	0.304 s

the outputs onto the mesh surface using nearest-neighbor or K -point averaging. For MVLM, input preparation includes rendering RGB views from 3D models; inference corresponds to heatmap-based landmark prediction; and final prediction involves back-projecting 2D predictions onto the 3D surface. For 2S-SGCN, input preparation includes graph construction; inference refers to the forward pass; and final prediction computes the MSE-over-mesh. PAL-Net achieves the lowest inference time, while its total runtime remains competitive despite additional patch creation and post-processing.

4.2. Anatomical analysis

In the following section, we present a comprehensive evaluation of PAL-Net on the LAFAS dataset based on two complementary analyses: point-wise and distance-wise, using results obtained through a 5-fold cross-validation framework. This evaluation strategy ensures that the reported performance reflects the model's ability to generalize across different subsets of the data while reducing bias introduced by a single train-test split.

4.2.1. Point-wise analysis

Euclidean distances between predicted and ground-truth landmarks were computed for each validation sample. Averaging these across folds yielded mean errors and standard deviations for all 50 landmarks. As detailed in Table 3, these metrics quantify the model's point-wise accuracy and generalization capability. Overall, PAL-Net achieved a mean localization error of 3.686 ± 2.306 mm, demonstrating robust and consistent performance across most landmarks. As expected, mid-line landmarks such as Subnasale (Sn), Stomion (Sto), and Nasion (N) exhibit relatively low errors, often below 2.5 mm. These landmarks tend to lie in well-defined, geometrically stable facial regions, which makes them easier to learn and to predict reliably. Similarly, paired landmarks on geometrically symmetrical areas such as Orbitales (Or), Cheek Points (Chk), and Labiale Superius/Inferius (Ls, Li) show closely matching errors between their left and right counterparts, indicating that PAL-Net maintains spatial symmetry and is not biased toward one side of the face. Landmarks located around the ear region, such as Gonion (Go), Tragion (T), Postaurale (Pa), and Superaurale (Sa), along with the Trichion (Tr), consistently exhibited the highest localization errors, frequently exceeding 5 mm and reaching over 7 mm in some cases. This elevated error is primarily attributed to corruption in the input mesh in these regions. For the majority of subjects, the 3D surface data around the ears and hairline were either incomplete, noisy, or poorly defined, largely due to occlusion and scattering introduced by hair. In the case of Trichion, the landmark was often located above the dense hairline, where the facial mesh abruptly degrades or ends, resulting in landmark positions that were not even on the mesh's valid surface. As a corrective measure during postprocessing, the mean position of the 10 closest vertices to the predicted landmark coordinate was assigned as the final predicted landmark to mitigate the effects of missing or corrupted geometry. Despite these efforts, these regions remain particularly challenging, and the high variability in surface quality limits the model's ability to consistently localize landmarks with high precision. The Gonion (Go) landmark consistently ranks among

the highest localization errors in 3D facial landmarking studies and was also among the least accurately predicted by PAL-Net. In the study by Aldridge et al. [40], which assessed the precision and measurement error of 3D soft tissue landmarks using photogrammetric systems, Gonion showed the highest localization error among all landmarks, with mean deviations up to 4.10 mm and a standard deviation of 1.64 mm. Moreover, inter-landmark distances involving Gonion were among those with the highest digitization errors, exceeding 5% of the total variance, indicating operator difficulty with consistent placement. This reduced accuracy is primarily due to the limited visibility of the landmark and the variable soft-tissue coverage in the mandibular angle region, making both manual and automated identification more challenging [41,42]. Interestingly, landmarks such as the Pogonion (Pg) and Crista Philtri (Cph) exhibit low mean errors (2.5–2.8 mm) despite being near highly curved regions, suggesting that the model generalizes well in structurally challenging areas when sufficient training data coverage is available. The standard deviation values across folds are generally low (0.297 mm on average), indicating stable model behavior and low sensitivity to fold-specific variability. This reflects both the anatomical consistency of the LAFAS dataset and the architectural regularization effect introduced by the attention modules in PAL-Net. Taken together, the results indicate that PAL-Net performs best in well-defined and midline regions, with consistent generalization across folds, while challenges remain in areas with surface noise or structural variability. These findings align with known limitations in 3D facial landmark localization and further emphasize the importance of geometric quality in input data for precise prediction.

To contextualize the performance of PAL-Net, it is important to compare the model's localization accuracy to known measures of intra-observer and inter-observer variability. In our case, the tight intra-observer variability across the dataset averages around 0.798 mm. PAL-Net achieves an overall mean localization error of 3.686 ± 2.306 mm across the 50 landmarks. While this error margin exceeds the precision baseline set by a single expert, it is critical to acknowledge that this performance is often constrained by data scarcity inherent to specialized medical datasets. Furthermore, the performance is not uniform, a substantial number of the 50 landmarks achieve significantly lower error, approaching human intra-observer consistency, while the overall mean is elevated due to higher variability in a smaller number of anatomically complex landmarks. Despite the precision gap with the gold standard, the 3.686 mm error remains within a clinically reasonable range for many applications. PAL-Net's value lies not in surpassing human precision, but in providing a consistent, standardized, and highly automated alternative.

4.2.2. Distance-wise analysis

To assess preservation of spatial relationships between landmarks, we performed a pairwise distance-variability analysis using 5-fold cross-validation. This involved computing inter-landmark distances and comparing their differences across the predicted and ground-truth datasets for each fold. The matrix shown in Fig. 4 is derived by calculating the average distance-wise error for each pair of landmarks across all test cases in the dataset. For each pair of landmarks, such as A and B, the Euclidean distance is computed independently for the ground-truth and predicted landmark positions. The absolute difference between these distances is then calculated for all test cases in the dataset, capturing how accurately the predicted landmarks maintain the spatial relationship between each pair. Finally, the differences are averaged across the entire set, yielding a symmetrical error matrix in which each entry represents the average absolute error for a specific pair of landmarks. The mean value of the matrix is 2.822 mm, representing the overall distance-wise error between the predicted and ground-truth landmarks. From the figure, we observe that the matrix shows areas where the model struggles to replicate ground-truth distances. Specifically, errors in the first row and column indicate that the model has difficulty preserving the spatial consistency of the

Table 3

Localization accuracy of anatomical landmarks, averaged over the validation sets of each fold, with standard deviation of the localization errors computed independently for each fold, then averaged across all folds.

Landmark	Intra- observer Reliability (mm)	Localization Accuracy (mm)
Trichion (Tr)	0.4289	6.413 ± 4.614
Glabella (G)	0.7238	2.672 ± 1.862
Nasion (N)	0.7766	2.240 ± 1.498
Pronasale (Prn)	0.4231	2.096 ± 1.194
Columella (C)	0.4932	1.824 ± 1.086
Subnasale (Sn)	0.4711	1.653 ± 0.967
Labiale Superius (Ls)	0.5104	2.035 ± 1.191
Stomion (Sto)	0.3407	1.640 ± 0.964
Labiale Inferius (Li)	0.5875	1.995 ± 1.251
Sublabiale (Sl)	0.9986	2.404 ± 1.558
Pogonion (Pg)	0.9200	2.750 ± 1.669
Gnathion (Gn)	0.8944	4.218 ± 3.279
Tragion Right (T)	0.6925	4.277 ± 2.161
Preaurale Right (Pra)	1.5077	5.437 ± 3.201
Superaurale Right (Sa)	1.3221	6.505 ± 3.965
Postaurale Right (Pa)	1.5267	6.940 ± 4.468
Subaurale Right (Sba)	1.2115	5.130 ± 2.903
Frontotemporale Right (Ft)	0.5389	4.823 ± 2.586
Zygion Right (Zy)	0.3994	5.345 ± 3.082
Gonion Right (Go)	0.3803	7.169 ± 3.916
Orbitale Superius Right (Os)	0.4817	3.407 ± 2.053
Exocanthion Right (Ex)	0.8051	2.806 ± 1.528
Orbitale Right (Or)	0.3592	5.613 ± 3.719
Endocanthion Right (En)	0.4633	2.206 ± 1.792
Malare Cheek Right (Chk)	1.5841	3.899 ± 2.314
Alar crest Right (Ac)	0.5654	2.202 ± 1.319
Alare Right (Al)	0.6430	1.960 ± 1.263
Inferior terminal point of the nostril axis Right (Itn)	0.8018	1.900 ± 1.064
Superior terminal point of the nostril axis Right (Stn)	0.5100	1.866 ± 1.057
Crista Philtri Right (Cph)	0.5949	2.134 ± 1.097
Cheilion Right (Ch)	0.4834	2.586 ± 1.496
Tragion Left (T)	0.7178	3.688 ± 2.011
Preaurale Left (Pra)	2.0102	5.059 ± 3.274
Superaurale Left (Sa)	1.5333	6.041 ± 3.902
Postaurale Left (Pa)	1.4705	6.709 ± 3.913
Subaurale Left (Sba)	1.7314	4.889 ± 2.813
Frontotemporale Left (Ft)	0.4314	4.629 ± 2.493
Zygion Left (Zy)	0.3389	5.619 ± 4.213
Gonion Left (Go)	0.4850	7.309 ± 5.048
Orbitale Superius Left (Os)	0.3231	3.576 ± 2.123
Exocanthion Left (Ex)	0.8059	3.226 ± 2.862
Orbitale Left (Or)	0.4385	5.779 ± 3.751
Endocanthion Left (En)	0.4930	2.491 ± 2.493
Malare Cheek Left (Chk)	2.1667	4.389 ± 2.864
Alar crest Left (Ac)	1.0431	2.204 ± 1.323
Alare Left (Al)	1.1591	2.068 ± 1.158
Inferior terminal point of the nostril axis Left (Itn)	0.7360	1.911 ± 1.121
Superior terminal point of the nostril axis Left (Stn)	0.5504	1.977 ± 1.094
Crista Philtri Left (Cph)	0.5784	2.105 ± 1.200
Cheilion Left (Ch)	0.4944	2.489 ± 1.534
Mean	0.798 mm	3.686 ± 2.306 mm

Trichion (Tr) landmark relative to all other landmarks. This suggests that the model’s predictions for Trichion are less reliable; similarly, the paired landmarks of Orbitale showed the largest error (the hottest point in the matrix). Conversely, most other landmark pairs exhibit smaller errors, showing that the model better maintains their spatial relationships.

Since this study focuses on anatomical landmarks for medical anthropometric analysis, not all linear distances on the face hold equal significance. Some facial linear distances are more distinct and clinically relevant than others [43]. To ensure meaningful analysis, specific anatomical linear distances and angles were selected for further evaluation. These measurements were then compared with the intra-observer variability to provide a comprehensive assessment of the predicted

Table 4

Mean absolute errors of selected linear anatomical distances, averaged over the validation sets of each fold. The table reports intra-observer variability, predicted distance errors, and the normalized error expressed as a percentage of the corresponding anatomical distance.

Linear distances	Intra variability (mm)	Error distances (mm)	Error distances Avg. distance (%)
Frontal Distances			
(Tr) — (N)	0.588	5.615	8.98%
(N) — (Pg)	0.890	2.726	2.64%
(N) — (Sn)	0.444	1.990	3.76%
(Sn) — (Pg)	0.840	2.290	4.43%
Horizontal Plane			
(Ex)R — (Ex)L	0.939	2.747	3.17%
(Zy)R — (Zy)L	0.081	1.942	1.47%
(T)R — (T)L	0.192	1.250	0.91%
(Ch)R — (Ch)L	0.652	2.709	5.68%
(Cph)R — (Cph)L	0.738	1.412	11.94%
(Go)R — (Go)L	0.328	4.101	3.69%
Sagittal Plane (Right)			
(T)R — (N)	0.459	2.321	2.02%
(T)R — (Sn)	0.415	2.454	2.07%
(T)R — (Pg)	0.544	2.703	2.04%
(Pg) — (Go)R	0.487	3.836	4.01%
(T)R — (Go)R	0.542	4.371	8.06%
Sagittal Plane (Left)			
(T)L — (N)	0.428	2.058	1.81%
(T)L — (Sn)	0.510	2.216	1.88%
(T)L — (Pg)	0.700	2.509	1.91%
(Pg) — (Go)L	0.384	3.975	4.18%
(T)L — (Go)L	0.503	4.172	7.70%
Mean	0.533 mm	2.870 mm	4.11%

facial landmarks and their accuracy. Table 4 demonstrates the intra-observer variability of anatomical distances, predicted landmarks’ distances, and the ratio between the predicted landmarks’ distances and the average distances of the landmarks. The ratio was calculated to provide insight into the significance of the error, depending on the distance. This approach allows evaluation of the magnitude of the error relative to the actual anatomical distances being measured, providing a clearer understanding of the model’s performance.

Despite the 2 mm benchmark being widely accepted as a clinically relevant threshold for facial landmark localization accuracy [44–46], not all linear distances predicted by our model fall within this margin. Anatomical distances vary greatly in scale, ranging from short inter-landmark distances such as (Cph)R—(Cph)L (11.8 mm, average distance) to broader measurements such as (Zy)R—(Zy)L (131.7 mm, average distance). In such cases, reporting error in absolute millimeters may not fully capture the model’s performance. Therefore, we normalized the error by the average ground-truth distances to provide relative insight. The majority of the predicted distances show errors well below 5% of the corresponding anatomical length, with a mean relative error of 4%. This ratio is informative when the absolute errors exceed 2 mm, while the proportional error remains within an acceptable range for clinical and anthropometric applications. When considered relative to landmark spacing and clinical practice tolerance, the results strongly support the feasibility of the proposed method for accurate anatomical landmark prediction. These results support the robustness of the model across both short and long facial spans, with predictive accuracy aligning well with clinical tolerances and intra-observer variability.

In addition to evaluating linear anatomical distances, we quantified how accurately PAL-Net reproduces clinically relevant anatomical angles defined by [47,48] (see Table 5). This table compares intra-observer variability and prediction accuracy for the anatomical angles, providing further insight into the accuracy and reliability of the predicted landmarks. Across nine standard angles, the mean absolute error was 1.28°, corresponding to just a 2.1% deviation from average anatomical values. Most midline and transverse angles, for




Fig. 4. Average distance-wise error matrix between ground truth and predicted landmarks averaged over the 5-fold cross validation. Each entry represents the absolute difference in pairwise distances for a specific landmark pair, averaged over the test set. The mean of the matrix is 2.822 mm, indicating the overall distance-wise error.

example, (T)R-(Prn)-(T)L and (N)-(Prn)-(Pg), showed errors below 1°, reflecting consistent performance in regions critical to facial symmetry. Larger deviations appeared in measures involving more complex mandibular contours, notably (Go)R-(Pg)-(Go)L (2.86°, 4% relative error) and (Sn)-(N)-(Prn) (1.44°, 8% relative error), likely due to inherent anatomical variability at these landmarks. Overall, PAL-Net’s angular accuracy closely approaches expert manual annotations and remains within accepted clinical tolerances for the majority of key anthropometric parameters, further supporting its use for automated, high-throughput 3D facial landmarking in both research and clinical workflows.

4.3. Ablation study

To assess the contributions of different aspects of the PAL-Net model, such as the attention mechanism and model depth, we conducted an ablation study involving alternative model configurations. Each variant was trained and evaluated using 5-fold cross-validation on the LAFAS dataset, maintaining identical preprocessing and training procedures for consistency. The first variant entirely removed the attention modules to evaluate their impact on landmark localization accuracy. The second employed a Top-K attention mechanism, selecting the *k* most informative points per attention layer to assess computational complexity and test whether selective context aggregation could improve performance. The third variant reduced the network depth by using two convolutional blocks instead of three, investigating the effect of a shallower model capacity on predictive accuracy. Table 6 summarizes the mean localization errors across folds for each configuration. None of the ablated models outperformed the baseline PAL-Net

Table 5

Mean absolute errors of selected anatomical angles, averaged over the validation sets of each fold. The table includes intra-observer variability, predicted angular errors, and the normalized error expressed as a percentage of the corresponding average anatomical angle.

Angular distances	Intra variability (°)	Error angles (°)	Error angle/ Avg. Angle (%)
(T)R - (N) - (T)L	0.095	1.190	1.33%
(T)R - (Prn) - (T)L	0.032	0.656	1.00%
(T)R - (Pg) - (T)L	0.345	0.446	0.66%
(Go)R - (Pg) - (Go)L	0.176	2.863	4.33%
(N) - (Sn) - (Pg)	1.435	1.260	1.00%
(N) - (Prn) - (Pg)	0.682	0.880	0.66%
(Sn) - (N) - (Prn)	0.228	1.440	8.00%
(T)R - (Go)R - (Pg)	1.618	1.916	1.66%
(T)L - (Go)L - (Pg)	0.676	0.910	1.00%
Mean	0.587°	1.284°	2.18%

with attention and three convolutional blocks. Specifically, the absence of attention led to a degradation in performance, highlighting the importance of global context modeling. Similarly, the Top-K attention variant did not yield improvements, resulting in higher complexity. The reduced-depth model also showed lower accuracy, indicating that sufficient network capacity is essential for capturing the complex spatial patterns required for precise 3D facial landmark localization. These results confirm that the original PAL-Net architecture balances model complexity and context integration effectively for this task.

Table 6

Ablation study results: mean localization and distance errors (in mm) averaged across 5-fold cross-validation.

Model Variant/Preprocessing	Localization error (mm)	Distance error (mm)
Preprocessing Variants (Applied to Baseline Architecture)		
Fixed-size patches (500 points)	3.731	2.866
Fixed-size patches (1500 points)	3.800	2.878
Radius-based patches (10 mm)	3.765	2.889
Radius-based patches (15 mm)	3.758	2.896
Radius-based patches (20 mm)	3.747	2.886
PAL-Net (coarse mesh)	3.985	3.041
PAL-Net (no spatial ordering)	3.773	2.872
PAL-Net (with RGB texture)	4.293	3.172
Model Architecture Variants		
Baseline PAL-Net (1000 points)	3.686	2.822
Without attention modules	3.713	2.851
Top-K attention ($k = 10$)	3.824	2.897
Reduced depth (2 conv layers)	3.909	2.964

Further analysis was performed to investigate the influence of preprocessing parameters on PAL-Net's landmark localization performance. The baseline model was trained using two distinct patch extraction strategies to evaluate its sensitivity to local neighborhood definitions and point cloud density. In the first approach, patches were extracted by selecting a fixed number of surface points closest to each approximated landmark, with patch sizes of $k = 500, 1000,$ and 1500 points. This strategy ensured consistent input dimensionality across samples, while allowing control over local detail granularity. In the second approach, patches were generated by including all surface points within fixed Euclidean radii of $D = 10$ mm, 15 mm, and 20 mm from the approximated landmark positions. This method simulates scenarios involving sparser or variable-resolution meshes, such as those commonly encountered in low-cost or mobile acquisition systems. The results demonstrate that the baseline PAL-Net model using 1000-point, fixed-size patches achieves the best overall performance, with localization and distance errors of 3.686 mm and 2.822 mm, respectively. Reducing the number of points to 500 slightly increases the error, while using 1500 points does not yield further improvement, suggesting a saturation point around 1000 points. Radius-based patching with radii of 10 – 20 mm results in comparable performance (3.74 – 3.76 mm localization error), indicating that increasing spatial extent does not provide significant benefit. These results emphasize that careful tuning of patch size, both in terms of point count and spatial coverage, is important to balance local detail capture and noise suppression. To evaluate PAL-Net's robustness to lower-resolution inputs, we tested its performance on coarsely sampled 3D facial meshes. Instead of constructing patches from the original high-resolution meshes, we downsampled each mesh to 10,000 points using uniform surface sampling. Local patches were then extracted from this coarse representation, selecting only the $K = 100$ nearest points around each approximated landmark, compared to $K = 1000$ in the baseline configuration. This setting simulates applications where acquisition devices produce sparse point clouds (e.g., mobile or low-cost systems). Despite the reduced point density and smaller patch size, PAL-Net retained performance, demonstrating its adaptability to varying mesh resolutions. To investigate the role of spatial ordering within each patch, we trained a variant of PAL-Net that does not enforce consistent point ordering. While point-wise convolutional layers are permutation-invariant, the fully connected layers at the end of the network aggregate information across all points, introducing sensitivity to input ordering. Without consistent ordering, the model receives arbitrary point arrangements for each patch, which hinders its ability to capture stable spatial patterns. Removing spatial ordering led to a measurable drop in performance, confirming that consistent point arrangement facilitates more effective feature aggregation and improves localization accuracy. The corresponding results are

Table 7

Comparison of point-wise and distance-wise localization errors (in mm) for different methods on the LAFAS dataset after excluding peripheral ear-region landmarks (8 out of 50 landmarks). Metrics reflect performance over the remaining 42 facial landmarks.

Model	Point-wise error (mm)	Distance-wise error (mm)
MVLM [19]	4.342	3.742
2S-SGCN [23]	3.196	3.076
PAL-Net (Ours)	3.276	2.570

reported in Table 6. To address the relevance of texture information, we evaluated a variant of PAL-Net that included RGB color values appended to the (x, y, z) coordinates of each point. As shown in Table 6, incorporating texture did not improve performance and actually led to a slight degradation in localization accuracy. This suggests that, for anatomical landmarking, geometric surface features are the primary discriminators, while texture data may introduce noise from lighting variations or skin tone that obscures the underlying morphological structure.

To evaluate whether prediction accuracy varied across facial regions, a Bland–Altman analysis was performed separately for midline, left, and right landmarks (Fig. 5). In all three regions, the mean difference between predicted and ground-truth coordinates was close to zero, indicating no systematic bias. Most differences fell within the 95% limits of agreement, with no strong trend suggesting under- or over-prediction in any specific region. The spread of differences was slightly wider for lateral landmarks, particularly on the left side, which may reflect localized variability in mesh quality or anatomical asymmetry. Overall, the plots suggest consistent model behavior across facial regions, with no evidence of directional error or regional bias.

Exclusion of peripheral landmarks

Certain landmarks located around the ears, namely *Preaurale (Pra)*, *Superaurale (Sa)*, *Postaurale (Pa)*, and *Subaurale (Sba)* on both the left and right sides, consistently exhibited higher localization errors across subjects. These points are often affected by mesh corruption caused by hair, leading to incomplete or noisy surface geometry. In several cases, the ground-truth landmark did not lie on the surface of any mesh point, making accurate prediction and evaluation inherently ambiguous. To better assess PAL-Net's performance and that of other methods in clinically relevant, geometrically stable regions, we excluded the eight peripheral landmarks and recomputed the localization and distance-based errors. Table 7 presents the updated evaluation, showing improved performance when focusing solely on core facial landmarks. This adjustment provides a more representative metric of the model's effectiveness for central facial analysis tasks.

On the reduced set of 42 facial landmarks (excluding the 8 ear-region points), 2S-SGCN has slightly lower point-wise accuracy than PAL-Net (3.196 mm vs. 3.276 mm). However, PAL-Net maintains the best performance in terms of distance-wise accuracy (2.570 mm), suggesting that its predictions better preserve anatomical structure. The distance-wise error of 2S-SGCN slightly increases compared to the full landmark set, indicating reduced spatial coherence when peripheral landmarks are excluded. This further underscores PAL-Net's strength in maintaining geometric consistency across core facial structures.

Discussion on model interpretability

Unlike traditional geometric methods that rely on explicit, hand-crafted features, deep learning models are often criticized for their lack of transparency. While we do not visualize internal feature maps, our experimental analysis provides critical insight into how PAL-Net arrives at its predictions by isolating the contribution of different spatial contexts. First, the ablation study (Table 6) explains the role of global context. The performance drop observed when removing the




Fig. 5. Bland–Altman plots showing prediction errors by facial region (midline, right, left). Each plot shows the difference between predicted and ground-truth coordinates relative to their mean. The dashed lines represent the mean difference (gray) and the 95% limits of agreement (red and blue). The results indicate no systematic bias across regions, with most predictions falling within acceptable error bounds.

attention mechanism confirms that local geometric features alone are insufficient for high-accuracy landmarking. The model explicitly relies on the attention-weighted aggregation of features from distant patches to resolve ambiguities. Second, the sensitivity analysis regarding patch size and radius explains the model’s reliance on local granularity. The results indicate a clear “structural sweet spot” (approximately 1000 points or 15–20mm radius). Performance degrades with smaller patches (insufficient context to orient the landmark) and with larger patches (introduction of irrelevant noise or background geometry). This demonstrates that PAL-Net’s decision-making is driven by a balanced integration of immediate local surface topology and the broader global arrangement of facial features. It is worth noting that while the quantitative differences between these configurations are relatively small, the consistent trend across cross-validation folds reinforces the need for both global attention and optimized local context to minimize outliers and achieve robust anatomical alignment.

5. Limitations, conclusions and future work

The proposed pipeline, integrating local patch-based points with global attention, demonstrated high accuracy, stability, and computational efficiency. PAL-Net consistently outperformed existing methods on both point-wise and distance-wise evaluations, achieving an average localization error of 3.686 mm on the LAFAS dataset, with a distance preservation error of 2.822 mm. This performance profile, achieved despite data scarcity supports the model’s feasibility as a high-throughput, reliable tool for 3D facial analysis. The lightweight nature of PAL-Net, requiring only 2.5 GiB of GPU memory and less than 4 s per training epoch, positions it as a viable solution for real-time clinical and research applications, including facial asymmetry analysis, pre- and post-operative planning, and growth assessment.

Although the proposed PAL-Net framework achieves high accuracy and computational efficiency in anatomical landmark localization, several limitations must be considered. The model was trained on a relatively small dataset (214 subjects) from a single laboratory (LAFAS), which may limit its generalizability to broader populations or pathological cases, despite 5-fold cross-validation. Performance also declined in regions with incomplete or noisy geometry, such as the ears, hairline, and jawline, due to occlusions and surface artifacts inherent to stereophotogrammetry, affecting both manual and automated annotations. Furthermore, the current pipeline is strictly dependent on the quality of the initial rigid registration of facial models to a shared reference frame. This requirement introduces a structural limitation, as misalignments or registration inaccuracies can negatively affect subsequent landmark approximation and prediction. As a result, the model lacks inherent invariance to rotation and translation. Another limitation lies in the initial landmark approximation based on averaged annotations from the training set. While effective for general structures, this introduces

bias in variable regions and may not generalize to atypical facial geometries. To address this, future work will extend the framework to pathological cases where landmarks may deviate significantly from normative patterns. Building a model capable of adapting to such variability is essential for expanding clinical applicability, particularly in cases involving craniofacial syndromes or surgical outcomes.

Future work will further focus on addressing the mentioned limitations by exploring model architectures that are intrinsically equivariant or invariant to rigid transformations. The integration of geometric learning techniques that handle partial, low-quality, or noisy data may further improve performance in challenging facial regions, such as corrupted points around the ears. Another promising direction for future research is to leverage the predicted soft-tissue landmarks to estimate corresponding skeletal landmarks visible in Computed Tomography (CT) scans [49]. Establishing reliable mappings between surface-based anatomical landmarks and their underlying skeletal counterparts could enable indirect estimation of craniofacial skeletal points without additional imaging. This approach has the potential to reduce patient exposure to ionizing radiation by minimizing reliance on CT solely for landmark localization, thereby contributing to safer, more efficient clinical workflows (see Fig. A.7).

CRedit authorship contribution statement

Ali Shadman Yazdi: Writing – original draft, Software, Methodology, Investigation, Formal analysis, Data curation. **Annalisa Cappella:** Writing – review & editing, Validation, Supervision, Funding acquisition, Data curation, Conceptualization. **Benedetta Baldini:** Writing – review & editing, Supervision, Software, Resources. **Riccardo Solazzo:** Formal analysis, Data curation. **Gianluca Tartaglia:** Supervision, Conceptualization. **Chiarella Sforza:** Supervision, Data curation, Conceptualization. **Giuseppe Baselli:** Visualization, Validation, Supervision, Software, Resources, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Special thanks to Marco Farronato for his helpful feedback and proofreading support during the writing of this manuscript.

Funding: The present research was partially funded by the University of Milan, Italy under the “My First SEED Grant” fund, DM 737/2021 MUR (Project: DIAERESSES - PSR_LINEA3 Piano di sviluppo di ricerca - Bando SoE-SEED- Linea 3).

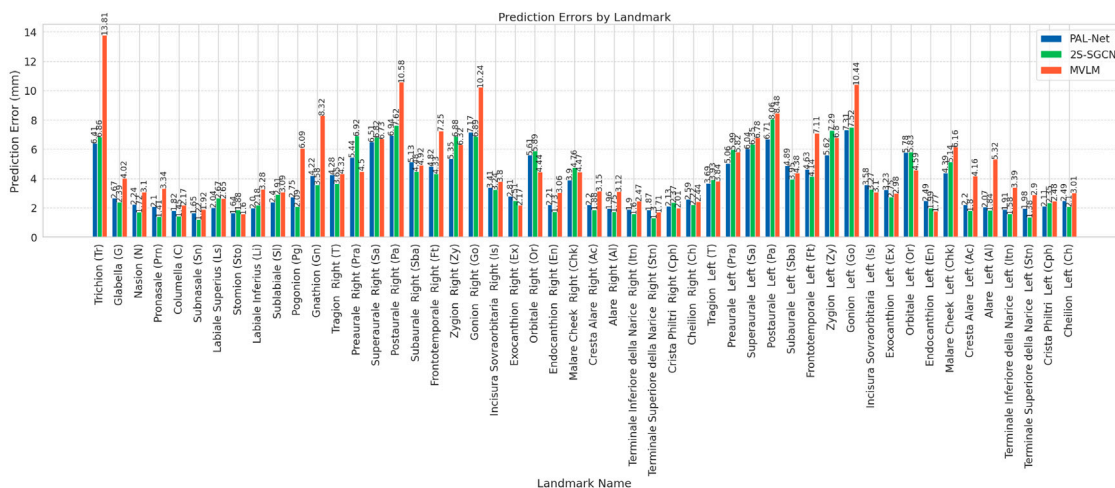


Fig. A.6. Localization error for each of the 50 landmarks in the LAFAS dataset, comparison of PAL-Net with existing methods.




Fig. A.7. Visualization of anatomical landmarks predicted by PAL-Net (red) compared to ground truth annotations (green) on a 3D facial mesh, shown from frontal and lateral views.

Appendix A. Per landmark model comparison

During the preparation of this work the author(s) used GPT 4o model in order to clarify and enhance the flow of the article for readers. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

References

[1] Masnada S, Gibelli D, Dolci C, De Giorgis V, Cappella A, Veggiotti P, Sforza C, Group IAS. 3D facial morphometry in Italian patients affected by icardi syndrome. *Am J Med Genet Part A* 2020;182(10):2325–32. <http://dx.doi.org/10.1002/ajmg.a.61791>, [Accessed 15 December 2025].

[2] Chang JB, Small KH, Choi M, Karp NS. Three-dimensional surface imaging in plastic surgery: foundation, practical applications, and beyond. *Plast Reconstr*

- Surg 2015;135(5):1295–304. <http://dx.doi.org/10.1097/prs.0000000000001221>, [Accessed 15 December 2025].
- [3] Codari M, Caffini M, Tartaglia GM, Sforza C, Baselli G. Computer-aided cephalometric landmark annotation for CBCT data. *Int J Comput Assist Radiol Surg* 2017;12:113–21. <http://dx.doi.org/10.1007/s11548-016-1453-9>, [Accessed 15 December 2025].
- [4] See MS, Roberts C, Nduka C. Age- and gravity-related changes in facial morphology: 3-dimensional analysis of facial morphology in mother-daughter pairs. *J Oral Maxillofac Surg* 2008-07-01;66(7):1410–6. <http://dx.doi.org/10.1016/j.joms.2007.12.041>, Publisher: Elsevier URL: [www.joms.org/article/S0278-2391\(08\)00090-6/fulltext](http://www.joms.org/article/S0278-2391(08)00090-6/fulltext).
- [5] Gibelli D, Cappella A, Bertozzi F, Sala D, Sitta S, Tasso DR, Tomasi F, Dolci C, Sforza C. Three-dimensional facial anthropometric analysis with and without landmark labelling: is there a real difference? *J Craniofac Surg* 2022;33(2):665–8. <http://dx.doi.org/10.1097/scs.00000000000007687>, [Accessed 15 December 2025].
- [6] Heike CL, Upson K, Stuhau E, Weinberg SM. 3D digital stereophotogrammetry: a practical guide to facial image acquisition. *Head Face Med* 2010-07-28;6(1):18. <http://dx.doi.org/10.1186/1746-160X-6-18>, [Accessed 15 December 2025].
- [7] Katina S, McNeil K, Ayoub A, Guilfoyle B, Khambay B, Siebert P, Sukno F, Rojas M, Vittert L, Waddington J, et al. The definitions of three-dimensional landmarks on the human face: an interdisciplinary view. *J Anat* 2016;228(3):355–65. <http://dx.doi.org/10.1111/joa.12407>, [Accessed 15 December 2025].
- [8] Facchi GM, Grossi G, D'Amelio A, Agnelli F, Sforza C, Tartaglia GM, Lanzarotti R. Graph neural networks for 3D facial morphology: Assessing the effectiveness of anthropometric and automated landmark detection. *Pattern Recognit Lett* 2025. <http://dx.doi.org/10.1016/j.patrec.2025.04.028>, [Accessed 15 December 2025].
- [9] Gibelli D, Dolci C, Cappella A, Sforza C. Reliability of optical devices for three-dimensional facial anatomy description: a systematic review and meta-analysis. *Int J Oral Maxillofac Surg* 2020;49(8):1092–106. <http://dx.doi.org/10.1016/j.ijom.2019.10.019>, [Accessed 15 December 2025].
- [10] Fagertun J, Harder S, Rosengren A, Moeller C, Werge T, Paulsen RR, Hansen TF. 3D facial landmarks: Inter-operator variability of manual annotation. *BMC Med Imaging* 2014;14:1–9. <http://dx.doi.org/10.1186/1471-2342-14-35>, [Accessed 15 December 2025].
- [11] Tang Y, Li H, Sun X, Morvan J-M, Chen L. Principal curvature measures estimation and application to 3D face recognition. *J Math Imaging Vision* 2017;59(2):211–33. <http://dx.doi.org/10.1007/s10851-017-0728-2>, [Accessed 15 December 2025].
- [12] Frantz S, Rohr K, Stiehl HS. Localization of 3D anatomical point landmarks in 3D tomographic images using deformable models. In: *International conference on medical image computing and computer-assisted intervention*. Springer; 2000, p. 492–501. http://dx.doi.org/10.1007/978-3-540-40899-4_50, [Accessed 15 December 2025].
- [13] Fan Z, Hu X, Chen C, Peng S. Boosting local shape matching for dense 3D face correspondence. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, p. 10944–54. <http://dx.doi.org/10.1109/cvpr.2019.01120>, [Accessed 15 December 2025].
- [14] Fan Z, Hu X, Chen C, Wang X, Peng S. A landmark-free approach for automatic, dense and robust correspondence of 3D faces. *Pattern Recognit* 2023;133:108971. <http://dx.doi.org/10.1016/j.patcog.2022.108971>, [Accessed 15 December 2025].
- [15] Creusot C, Pears N, Austin J. A machine-learning approach to keypoint detection and landmarking on 3D meshes. *Int J Comput Vis* 2013;102(1):146–79. <http://dx.doi.org/10.1007/s11263-012-0605-9>, [Accessed 15 December 2025].
- [16] O'Sullivan E. Extending convolutional pose machines for facial landmark localization in 3D point clouds. In: *Proceedings of the IEEE/CVF international conference on computer vision workshops*. 2019, <http://dx.doi.org/10.1109/iccvw.2019.00564>, [Accessed 15 December 2025].
- [17] Wei S-E, Ramakrishna V, Kanade T, Sheikh Y. Convolutional pose machines. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, p. 4724–32. <http://dx.doi.org/10.1109/cvpr.2016.511>, [Accessed 15 December 2025].
- [18] Qi CR, Yi L, Su H, Guibas LJ. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Adv Neural Inf Process Syst* 2017;30. <http://dx.doi.org/10.1109/cvpr.2017.16>, [Accessed 15 December 2025].
- [19] Paulsen RR, Juhl KA, Haspang TM, Hansen T, Ganz M, Einarsson G. Multi-view consensus CNN for 3D facial landmark placement. In: *Jawahar CV, Li H, Mori G, Schindler K, editors. Computer vision – ACCV 2018. Lecture notes in computer science*, Springer International Publishing; 2019, p. 706–19. http://dx.doi.org/10.1007/978-3-030-20887-5_44, [Accessed 15 December 2025].
- [20] Yin L, Wei X, Sun Y, Wang J, Rosato MJ. A 3D facial expression database for facial behavior research. In: *7th international conference on automatic face and gesture recognition*. IEEE; 2006, p. 211–6. <http://dx.doi.org/10.1109/fgcr.2006.6>, [Accessed 15 December 2025].
- [21] Wang Y, Cao M, Fan Z, Peng S. Learning to detect 3D facial landmarks via heatmap regression with graph convolutional network. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 36, 2022, p. 2595–603. <http://dx.doi.org/10.1609/aaai.v36i3.20161>, [Accessed 15 December 2025].
- [22] Xu M, Ding R, Zhao H, Qi X. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, p. 3173–82. <http://dx.doi.org/10.1109/cvpr46437.2021.00319>, [Accessed 15 December 2025].
- [23] Burger J, Blandano G, Facchi GM, Lanzarotti R. 2S-SGCN: A two-stage stratified graph convolutional network model for facial landmark detection on 3D data. *Comput Vis Image Underst* 2024;104227. <http://dx.doi.org/10.1016/j.cviu.2024.104227>, [Accessed 15 December 2025].
- [24] Yang H, Zhu H, Wang Y, Huang M, Shen Q, Yang R, Cao X. FaceScape: a large-scale high quality 3D face dataset and detailed riggable 3D face prediction. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2020, <http://dx.doi.org/10.1109/cvpr42600.2020.00068>, [Accessed 15 December 2025].
- [25] Zhang J, Liu M, Shen D. Detecting anatomical landmarks from limited medical imaging data using two-stage task-oriented deep neural networks. *IEEE Trans Image Process* 2017;26(10):4753–64. <http://dx.doi.org/10.1109/tip.2017.2721106>, [Accessed 15 December 2025].
- [26] Chong Y, Du F, Ma X, An Y, Huang Q, Long X, Huang J, Li Z, Yu N, Wang X. Automated anatomical landmark detection on 3D facial images using U-NET-based deep learning algorithm. *Quant Imaging Med Surg* 2024;14(3):2466. <http://dx.doi.org/10.21037/qims-22-1108>, [Accessed 15 December 2025].
- [27] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. In: *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer; 2015, p. 234–41. http://dx.doi.org/10.1007/978-3-662-54345-0_3, [Accessed 15 December 2025].
- [28] Guo J, Mei X, Tang K. Automatic landmark annotation and dense correspondence registration for 3D human facial images. *BMC Bioinformatics* 2013;14:1–12. <http://dx.doi.org/10.1186/1471-2105-14-232>, [Accessed 15 December 2025].
- [29] Berends B, Bielevelt F, Schreurs R, Vinayahalingam S, Maal T, de Jong G. Fully automated landmarking and facial segmentation on 3D photographs. *Sci Rep* 2024;14(1):6463. <http://dx.doi.org/10.1038/s41598-024-56956-9>, [Accessed 15 December 2025].
- [30] Zhang F, Mao S, Li Q, Peng X. 3D landmark detection on human point clouds: A benchmark and a dual cascade point transformer framework. *Expert Syst Appl* 2026;301:130425. <http://dx.doi.org/10.1016/j.eswa.2025.130425>, [Accessed 15 December 2025].
- [31] Kubík T, Kodým O, Šilling P, Trávníčková K, Mojžiš T, Matula J. Leveraging point transformers for detecting anatomical landmarks in digital dentistry. In: *International conference on medical image computing and computer-assisted intervention*. Springer; 2024, p. 216–28. http://dx.doi.org/10.1007/978-3-031-88977-6_20, [Accessed 15 December 2025].
- [32] Manal ER, Arsalane Z, Aicha M. Survey on the approaches based geometric information for 3D face landmarks detection. *IET Image Process* 2019;13(8):1225–31. <http://dx.doi.org/10.1049/iet-ipr.2018.6117>, [Accessed 15 December 2025].
- [33] Hua B-S, Tran M-K, Yeung S-K. Pointwise convolutional neural networks. 2018, p. 984–93. <http://dx.doi.org/10.1109/cvpr.2018.00109>, https://openaccess.thecvf.com/content_cvpr_2018/html/Hua_Pointwise_Convolutional_Neural_CVPR_2018_paper.html, [Accessed 15 December 2025].
- [34] De Stefani A, Barone M, Hatami Alamdari S, Barjami A, Baciliero U, Apolloni F, Gracco A, Bruno G. Validation of vectra 3D imaging systems: a review. *Int J Environ Res Public Health* 2022;19(14):8820. <http://dx.doi.org/10.3390/ijerph19148820>, [Accessed 15 December 2025].
- [35] Ferrario VF, Sforza C, Serrao G, Ciusa V, Dellavia C. Growth and aging of facial soft tissues: A computerized three-dimensional mesh diagram analysis. *Clin Anat: Off J Am Assoc Clin Anat Br Assoc Clin Anat* 2003;16(5):420–33. <http://dx.doi.org/10.1002/ca.10154>, [Accessed 15 December 2025].
- [36] Cappella A, Solazzo R, Yang J, Hassan NM, Dolci C, Gibelli D, Tartaglia G, Sforza C. Facial asymmetry of Italian children: a cross-sectional analysis of three-dimensional stereophotogrammetric reference values. *Symmetry* 2023;15(4):792. <http://dx.doi.org/10.3390/sym15040792>, [Accessed 15 December 2025].
- [37] Solazzo R, Cappella A, Gibelli D, Dolci C, Tartaglia G, Sforza C. Three-dimensional geometric morphometric characterization of facial sexual dimorphism in juveniles. *Diagnostics* 2025;15(3):395. <http://dx.doi.org/10.3390/diagnostics15030395>, [Accessed 15 December 2025].
- [38] Zhu H, Yang H, Guo L, Zhang Y, Wang Y, Huang M, Wu M, Shen Q, Yang R, Cao X. FaceScape: 3D facial dataset and benchmark for single-view 3D face reconstruction. *IEEE Trans Pattern Anal Mach Intell (TPAMI)* 2023. <http://dx.doi.org/10.1109/tpami.2023.3307338>, [Accessed 15 December 2025].
- [39] Sforza C, Dellavia C, De Menezes M, Rosati R, Ferrario VF. Three-dimensional facial morphometry: from anthropometry to digital morphology. In: *Handbook of anthropometry: physical measures of human form in health and disease*. Springer; 2012, p. 611–24. http://dx.doi.org/10.1007/978-1-4419-1788-1_32, [Accessed 15 December 2025].
- [40] Aldridge K, Boyadjiev SA, Capone GT, DeLeon VB, Richtsmeier JT. Precision and error of three-dimensional phenotypic measures acquired from 3dMD photogrammetric images. *Am J Med Genet Part A* 2005;138(3):247–53. <http://dx.doi.org/10.1002/ajmg.a.30959>, [Accessed 15 December 2025].

- [41] Staller S, Anigbo J, Stewart K, Dutra V, Turkkahraman H. Precision and accuracy assessment of single and multicamera three-dimensional photogrammetry compared with direct anthropometry. *Angle Orthod* 2022;92(5):635–41. <http://dx.doi.org/10.2319/101321-770.1>, [Accessed 15 December 2025].
- [42] Nord F, Ferjencik R, Seifert B, Lanzer M, Gander T, Matthews F, Rücker M, Lübbers H-T. The 3dMD photogrammetric photo system in cranio-maxillofacial surgery: Validation of interexaminer variations and perceptions. *J Cranio-Maxillofacial Surg* 2015;43(9):1798–803. <http://dx.doi.org/10.1016/j.jcms.2015.08.017>, [Accessed 15 December 2025].
- [43] Pucciarelli V, Bertoli S, Codari M, De Amicis R, De Giorgis V, Battezzati A, Veggiotti P, Sforza C. The face of Glut1-DS patients: A 3D craniofacial morphometric analysis. *Clin Anat* 2017;30(5):644–52. <http://dx.doi.org/10.1002/ca.22890>, [Accessed 15 December 2025].
- [44] Dindaroğlu F, Kutlu P, Duran GS, Görgülü S, Aslan E. Accuracy and reliability of 3D stereophotogrammetry: a comparison to direct anthropometry and 2D photogrammetry. *Angle Orthod* 2016;86(3):487–94. <http://dx.doi.org/10.2319/041415-244.1>, [Accessed 15 December 2025].
- [45] Othman SA, Saffai L, Wan Hassan WN. Validity and reproducibility of the 3D VECTRA photogrammetric surface imaging system for the maxillofacial anthropometric measurement on cleft patients. *Clin Oral Investig* 2020;24:2853–66. <http://dx.doi.org/10.1007/s00784-019-03150-1>, [Accessed 15 December 2025].
- [46] Weinberg SM, Naidoo S, Govier DP, Martin RA, Kane AA, Marazita ML. Anthropometric precision and accuracy of digital three-dimensional photogrammetry: comparing the genex and 3dmd imaging systems with one another and with direct anthropometry. *J Craniofacial Surg* 2006;17(3):477–83. <http://dx.doi.org/10.1597/03-066.1>, [Accessed 15 December 2025].
- [47] de Menezes M, Rosati R, Ferrario VF, Sforza C. Accuracy and reproducibility of a 3-dimensional stereophotogrammetric imaging system. *J Oral Maxillofac Surg* 2010;68(9):2129–35. <http://dx.doi.org/10.1016/j.joms.2009.09.036>, [Accessed 15 December 2025].
- [48] Gibelli D, Pucciarelli V, Cappella A, Dolci C, Sforza C. Are portable stereophotogrammetric devices reliable in facial imaging? a validation study of VECTRA H1 device. *J Oral Maxillofacial Surg: Off J Am Assoc Oral Maxillofacial Surg* 2018;76(8):1772–84. <http://dx.doi.org/10.1016/j.joms.2018.01.021>, [Accessed 15 December 2025].
- [49] Serafin M, Baldini B, Cabitza F, Carrafiello G, Baselli G, Del Fabbro M, Sforza C, Caprioglio A, Tartaglia GM. Accuracy of automated 3D cephalometric landmarks by deep learning algorithms: systematic review and meta-analysis. *La Radiol Medica* 2023;128(5):544–55. <http://dx.doi.org/10.1007/s11547-023-01629-2>, [Accessed 15 December 2025].