

NICE: Neural Implicit Craniofacial Model for Orthognathic Surgery Prediction

Jiawen Yang^{1*}, Yihui Cao^{1*}, Xuanyu Tian¹, Yuyao Zhang¹, Hongjiang Wei^{2†}

¹ School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China

² School of Biomedical Engineering, Shanghai Jiao Tong University, Shanghai 200127, China
yangjw2023@shanghaitech.edu.cn, caoyh2025@shanghaitech.edu.cn, tianxy@shanghaitech.edu.cn,
zhangyy8@shanghaitech.edu.cn, hongjiang.wei@sjtu.edu.cn

Abstract

Orthognathic surgery is a crucial intervention for correcting dentofacial skeletal deformities to enhance occlusal functionality and facial aesthetics. Accurate postoperative facial appearance prediction remains challenging due to the complex nonlinear interactions between skeletal movements and facial soft tissue. Existing biomechanical, parametric models and deep-learning approaches either lack computational efficiency or fail to fully capture these intricate interactions. To address these limitations, we propose Neural Implicit Craniofacial Model (NICE) which employs implicit neural representations for accurate anatomical reconstruction and surgical outcome prediction. NICE comprises a shape module, which employs region-specific implicit Signed Distance Function (SDF) decoders to reconstruct the facial surface, maxilla, and mandible, and a surgery module, which employs region-specific deformation decoders. These deformation decoders are driven by a shared surgical latent code to effectively model the complex, nonlinear biomechanical response of the facial surface to skeletal movements, incorporating anatomical prior knowledge. The deformation decoders output point-wise displacement fields, enabling precise modeling of surgical outcomes. Extensive experiments demonstrate that NICE outperforms current state-of-the-art methods, notably improving prediction accuracy in critical facial regions such as lips and chin, while robustly preserving anatomical integrity. This work provides a clinically viable tool for enhanced surgical planning and patient consultation in orthognathic procedures.

Code — <https://github.com/Bill-Yang-1/NICE>

1 Introduction

Orthognathic surgery is a principal clinical intervention for maxillofacial deformities, undertaken to correct dentofacial skeletal malocclusions and thereby improve both occlusal functionality and facial aesthetics (Posnick 2021). Within the surgical workflow, preoperative planning is especially critical: accurately forecasting postoperative soft-tissue appearance not only facilitates surgeon–patient communication and optimizes surgical strategy but also improves over-

all patient satisfaction (Stokbro et al. 2014). Facial soft tissue responds to skeletal movement through highly complex, nonlinear couplings that remain only partially understood, making postoperative facial appearance prediction a formidable challenge.

Early solutions focused on biomechanical model notably the mass-spring models (MSMs), mass-tensor models (MTMs) and finite-element models (FEMs) (Nadjmi et al. 2014; Knoops et al. 2018; Kim et al. 2019). An incremental FEM with realistic lip-sliding effects (FEM-RLSE) currently delivers the most precise perioral predictions, yet its accuracy depends on labour-intensive, patient-specific model construction, resulting in lengthy preparation times and high computational cost (Kim et al. 2021). These factors restrict its scalability and real-time applicability in routine clinical settings.

Three-dimensional parametric facial models have advanced facial geometry reconstruction by introducing structural constraints. Models such as Faces Learned with an Articulated Model and Expressions (FLAME) (Li et al. 2017) and Neural Parametric Head Models (NPHM) (Giebenhain et al. 2023) achieve efficient soft tissue fitting but these models only fit the soft tissue, omitting the mandible and maxilla. SCULPTOR (Qiu et al. 2022) extends this paradigm by jointly modelling the skull and facial soft tissue with a skeleton-driven strategy that maintains anatomical consistency and enhances realism, showing promise for postoperative facial appearance prediction. Nevertheless, its reliance on a linear blend skinning (LBS) limits its ability to capture the complex nonlinear deformation between skeleton and soft tissue.

The rapid progress of deep learning offers a new route for postoperative facial appearance prediction. Data-driven networks (Lampen et al. 2023; Ma et al. 2022, 2023; Fang et al. 2024; Bao et al. 2024; Huang et al. 2025) learn the intricate nonlinear mapping between skeletal movements and soft-tissue deformation from paired preoperative and postoperative data, markedly improving computational efficiency and automation. Building on this paradigm, several studies (Yang et al. 2025) have augmented the SCULPTOR framework with neural components, capturing subtle interactions between skeleton and soft tissue. Nevertheless, these approaches—like most current neural models—are supervised mainly by point-to-point distances. This neglect of

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

explicit geometric relationships among neighbouring points and surface connectivity hampers continuity and normal-vector consistency. Moreover, the absence of anatomical and kinematic constraints can allow predicted skeletal meshes to penetrate the skin surface, compromising structural fidelity and clinical reliability.

To address these limitations while simultaneously ensuring anatomical consistency and expressive nonlinear modeling, we propose Neural Implicit Craniofacial Model (NICE) for orthognathic surgery prediction. Our key contributions are summarized as follows:

- A shape module employing region-specific implicit SDF decoders (for face, maxilla, mandible) built around localized landmarks, enabling high-fidelity reconstruction of both bony structures and facial surface.
- A surgery module utilizing region-specific deformation decoders driven by a shared surgery latent code, effectively modeling the complex, nonlinear biomechanical response of facial surface to skeletal movements while leveraging anatomical priors.
- An integrated framework providing a streamlined workflow for accurate postoperative facial appearance prediction. Extensive experiments demonstrate that NICE outperforms state-of-the-art methods in prediction fidelity, particularly in critical regions like lips and chin, while robustly maintaining anatomical consistency.

2 Related Work

2.1 Conventional Biomechanical Models

Conventional biomechanical approaches to postoperative facial appearance prediction fall into three main categories: the mass–spring models (MSMs), mass–tensor models (MTMs) and finite element methods (FEMs) (Nadjmi et al. 2014; Knoops et al. 2018; Kim et al. 2019). MSM models soft tissue as point masses connected by springs governed by Hooke’s law. MTM extends MSM by assigning each mass a local stiffness tensor, enabling spatially localized, direction-dependent rigidity and hence anisotropic stress–strain behavior. FEM discretizes maxillofacial skeleton and soft tissue into three-dimensional elements and solves linear-elastic or nonlinear equations. The incremental FEM with realistic lip-sliding effects (FEM-RLSE) imposes explicit lip-sliding constraints and staged skeletal increments, which is widely regarded as the most precise simulator of postoperative facial change (Kim et al. 2021). Despite their accuracy, biomechanical models entail patient-specific meshing and heavy computation. End-to-end processing times frequently reach several hours per case, severely limiting scalability and real-time applicability in routine clinical practice.

2.2 Head Parametric Models

Parametric models encode craniofacial variation in a differentiable, low-dimensional latent space, offering an attractive compromise between computational efficiency and controllability. FLAME (Li et al. 2017) learns identity and expression shape spaces via PCA and employs LBS

to couple them. NPHM (Giebenhain et al. 2023) represents facial geometry as a continuous SDF, achieving resolution-independent detail. However, most parametric models deform only the soft tissue surface, omitting skeleton-driven mechanics essential to orthognathic planning. To address this gap, SCULPTOR (Qiu et al. 2022) extends the human face model FLAME (Li et al. 2017) to a physically precise parametric human skeleton model with the canonical pose. It introduces a trait component to model physiologically plausible variations in the internal skeletal structure of the face, as well as the associated shape changes that occur due to modifications in the skeleton. The trait component improves the model’s ability to generate more characteristic faces with a physiologically correct constraint. Nevertheless, the linear additive nature of LBS limits its ability to capture nonlinear relationships. This limitation constrains SCULPTOR and similar parametric models from fully representing the intricate nonlinear interactions between craniofacial skeleton and soft tissue. Thus they fall short of the precision demanded by clinical postoperative facial appearance prediction.

2.3 Deep Learning Models

Recent deep learning advances have led to more efficient methods for predicting postoperative facial appearance. PointNet (Qi et al. 2017) introduced end-to-end geometric-feature learning on unordered point clouds, eliminating the need for meshes or voxelized representations. Inspired by this, subsequent studies represent skull–face morphology as three-dimensional point clouds and train networks to learn per-point displacement fields that map preoperative to postoperative anatomy. A representative example is ACMT-Net (Fang et al. 2024) which introduces an Attention Correspondence assisted Movement Transformation network to correlate skeletal displacements with facial soft-tissue changes, achieving higher computational efficiency than FEM-based methods while maintaining accuracy. Another approach (Yang et al. 2025) integrates neural networks into the SCULPTOR parametric framework, modeling nonlinear skeleton-to-soft-tissue deformation and improving reconstruction accuracy. However, most networks use Chamfer distance or point-to-point losses without explicit constraints on local geometry and mesh connectivity. This can lead to discontinuities, inconsistent normals, and physiologically implausible deformations, limiting the reliability needed for clinical use.

3 Method

3.1 Overview

Our neural implicit craniofacial model primarily consist of two modules: the shape module and the surgery module. The shape module represents the anatomical geometry of the facial outer surface and underlying skull within the canonical space. The core concept of the surgery module centers on biomechanical deformations, which constrain the plausible transformation space for both the facial outer surface and skull during surgical procedures. In addition, we present

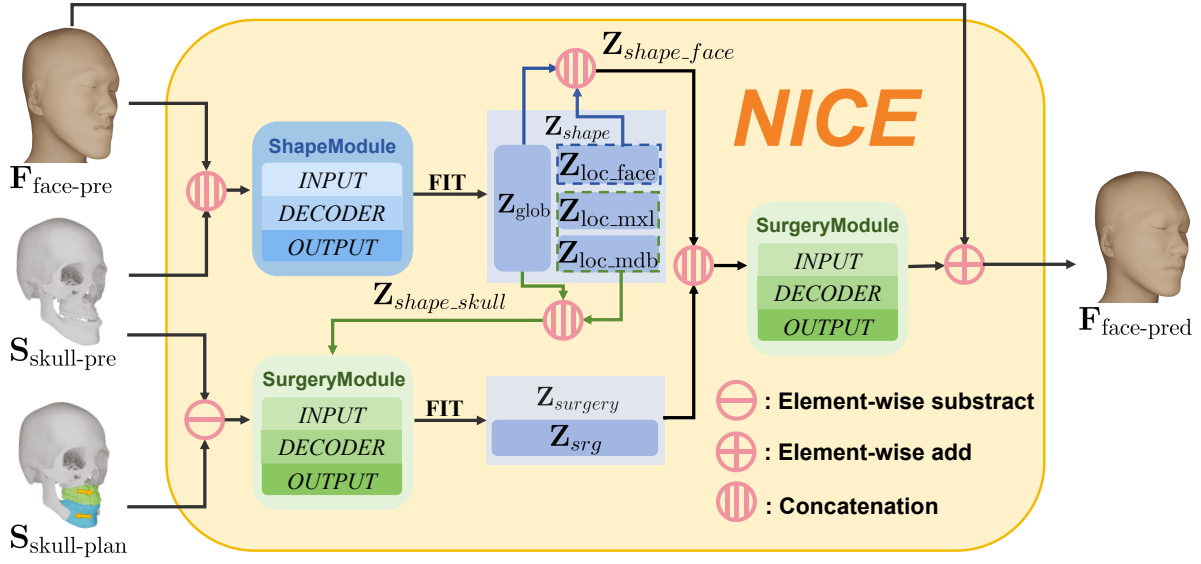


Figure 1: Pipeline of our neural implicit craniofacial model (NICE) for postoperative facial appearance prediction. The shape module first reconstructs the preoperative facial surface $\mathbf{F}_{\text{face-pre}}$ and skull $\mathbf{S}_{\text{skull-pre}}$ from CT scans, generating the corresponding shape latent codes $\mathbf{Z}_{\text{shape_face}}$ and $\mathbf{Z}_{\text{shape_skull}}$. The surgery module then fits a shared surgery latent code \mathbf{Z}_{srg} by matching the deformation from $\mathbf{S}_{\text{skull-pre}}$ to the surgically planned postoperative skull $\mathbf{S}_{\text{skull-plan}}$. Finally, \mathbf{Z}_{srg} , $\mathbf{Z}_{\text{shape_face}}$ and $\mathbf{F}_{\text{face-pre}}$ are combined in the surgery module to predict the postoperative facial surface $\mathbf{F}_{\text{face-pred}}$.




Figure 2: Architecture of the shape module in NICE. Region-specific SDF decoders reconstruct facial, maxillary, and mandibular geometry from spatial coordinates and shape latent codes, using landmark-centered MLP ensembles.

a postoperative facial appearance prediction framework as shown in Fig. 1.

3.2 Shape Module

We implicitly represent person-specific head geometry using canonical-space signed distance functions (SDFs). As shown in Fig. 2, our shape module employs three region-specific SDF decoders for the facial surface, maxilla, and mandible. For each SDF decoder, we decompose it into an ensemble of several smaller MLPs, each constructed around K specific landmarks $\mathbf{lmk} \in \mathbf{R}^{K \times 3}$. We define the inte-

grated SDF decoder through the following equation:

$$\mathcal{F}_{\text{shape}}(\mathbf{x}, \mathbf{Z}_{\text{shape}}) = \sum_{k=0}^K w_k(\mathbf{x}, \mathbf{lmk}^k) f_{\theta}^k(\mathbf{x}, [\mathbf{Z}_{\text{glob}}, \mathbf{Z}_{\text{loc}}^k]). \quad (1)$$

Specifically, the inputs to each SDF decoder are: (1) query spatial coordinates $\mathbf{x} \in \mathbf{R}^3$, and (2) region-specific shape latent code $\mathbf{Z}_{\text{shape}}$ which represents the geometric characteristics of its anatomical region: face, maxilla or mandible.

$\mathbf{Z}_{\text{shape}}$ consists of a globally shared latent vector $\mathbf{Z}_{\text{glob}} \in \mathbf{R}^{d_{\text{glob}}}$ and K region-specific local latent vectors $\mathbf{Z}_{\text{loc}}^k \in \mathbf{R}^{d_{\text{loc}}}$. The same \mathbf{Z}_{glob} is used across the facial, maxillary, and mandibular SDF decoders, while each $\mathbf{Z}_{\text{loc}}^k$ corresponds to the k -th latent vector from K localized MLPs, associated with the facial surface (denoted as $\mathbf{Z}_{\text{loc_face}}$), maxilla ($\mathbf{Z}_{\text{loc_mxl}}$), or mandible ($\mathbf{Z}_{\text{loc_mdb}}$), respectively. To enhance \mathbf{Z}_{glob} 's representation capacity and constrain the semantic scope of local MLPs, we employ a compact MLP_{pos} to predict these K landmark positions \mathbf{lmk} from \mathbf{Z}_{glob} . The weights $w_k(\mathbf{x}, \mathbf{lmk}^k)$ in Eq. (1) are computed using Gaussian kernels, as commonly done in (Genova et al. 2020; Zheng et al. 2022). The f_{θ}^k in Eq. (1) denotes the k -th small MLP $_{\theta}$, defined as follows:

$$f_{\theta}^k(\mathbf{x}, [\mathbf{Z}_{\text{glob}}, \mathbf{Z}_{\text{loc}}^k]) = \text{MLP}_{\theta}^k(\mathbf{x} - \mathbf{lmk}^k, [\mathbf{Z}_{\text{glob}}, \mathbf{Z}_{\text{loc}}^k]), \quad (2)$$

where $[\cdot]$ represents the concatenation operation. Specifically, to handle SDF values in regions distant from all landmarks in \mathbf{lmk} , we set $\mathbf{lmk}^0 = \mathbf{0} \in \mathbf{R}^3$ and introduce the following definition:

$$f_{\theta}^0(\mathbf{x}, [\mathbf{Z}_{\text{glob}}, \mathbf{Z}_{\text{loc}}^0]) = \text{MLP}_{\theta}^0(\mathbf{x}, [\mathbf{Z}_{\text{glob}}, \mathbf{Z}_{\text{loc}}^0]). \quad (3)$$

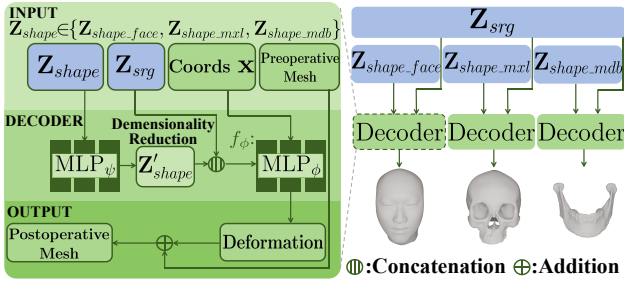


Figure 3: Architecture of the surgery module in NICE. Region-specific decoders predict deformation fields from spatial coordinates, a shared surgery latent code, and shape latent priors, enabling personalized postoperative mesh generation.

Since the SDF decoder outputs SDF values at query coordinates \mathbf{x} , we can then extract the region-specific mesh through spatial sampling via the marching cubes algorithm.

3.3 Surgery Module

The surgery module models the transformation of the facial surface and skull from preoperative to postoperative states using deformation fields. As shown in Fig. 3, it incorporates three region-specific decoders for the facial surface, maxilla, and mandible, similar to the shape module. We define this decoder through the following equation:

$$\mathcal{F}_{srg}(\mathbf{x}, \mathbf{Z}_{srg}, \mathbf{Z}_{shape}) = f_{\phi}(\mathbf{x}, [\mathbf{Z}_{srg}, \text{MLP}_{\psi}(\mathbf{Z}_{glob}, \mathbf{Z}_{loc})]). \quad (4)$$

Here, f_{ϕ} denotes the main MLP_{ϕ} that predicts the deformation field based on the concatenated latent representation. Specifically, the inputs to each decoder are: (1) query spatial coordinates \mathbf{x} , (2) surgery latent code $\mathbf{Z}_{srg} \in \mathbf{R}^{d_{srg}}$ simultaneously constraining the plausible transformation space for facial surface and skull structures during surgical procedures, and (3) shape latent code \mathbf{Z}_{shape} , serving as anatomical prior, obtained from the shape module by fitting the geometry of the corresponding region. The surgery latent code \mathbf{Z}_{srg} is globally shared across the three decoders, i.e., the same \mathbf{Z}_{srg} is used for the facial surface, maxilla, and mandible deformation. The MLP_{ψ} in Eq. (4) performs dimensionality reduction on \mathbf{Z}_{shape} to reduce computational cost. Since this decoder outputs the displacement from preoperative to postoperative states at query coordinates \mathbf{x} , we can obtain the postoperative mesh by applying this displacement to the preoperative person-specific head mesh.

3.4 Postoperative Facial Appearance Prediction

The pipeline of our postoperative facial appearance prediction framework is shown in Fig. 1. According to Eq. (1), we first employ the shape module to reconstruct the patient’s preoperative facial surface $\mathbf{F}_{\text{face-pre}}$, and skull structures maxilla $\mathbf{S}_{\text{mxl-pre}}$ and mandible $\mathbf{S}_{\text{mdb-pre}}$ obtained from CT scan data:

$$\begin{aligned} \mathbf{F}_{\text{face-pre}} &= \mathcal{F}_{shape}(\mathbf{x}_{\text{face}}, \mathbf{Z}_{\text{shape-face}}) \\ \mathbf{S}_{\text{mxl-pre}} &= \mathcal{F}_{shape}(\mathbf{x}_{\text{mxl}}, \mathbf{Z}_{\text{shape-mxl}}) \\ \mathbf{S}_{\text{mdb-pre}} &= \mathcal{F}_{shape}(\mathbf{x}_{\text{mdb}}, \mathbf{Z}_{\text{shape-mdb}}) \end{aligned} \quad (5)$$

This yields patient-specific shape latent codes $\mathbf{Z}_{\text{shape-face}}$, $\mathbf{Z}_{\text{shape-mxl}}$, and $\mathbf{Z}_{\text{shape-mdb}}$, encoding the unique geometry of each anatomical region. For brevity, we denote the combined skull structures and their corresponding shape latent codes as $\mathbf{S}_{\text{skull-pre}} = \{\mathbf{S}_{\text{mxl-pre}}, \mathbf{S}_{\text{mdb-pre}}\}$ and $\mathbf{Z}_{\text{shape-skull}} = \{\mathbf{Z}_{\text{shape-mxl}}, \mathbf{Z}_{\text{shape-mdb}}\}$, respectively.

Based on Eq. (4), we use the surgery module to fit the shared surgery latent code \mathbf{Z}_{srg} to represent the planned bony movements. This is achieved by matching the deformation between the preoperative skull $\mathbf{S}_{\text{skull-pre}}$ and the surgically planned postoperative skull $\mathbf{S}_{\text{skull-plan}}$:

$$\mathcal{F}_{srg}(\mathbf{x}_{\text{skull}}, \mathbf{Z}_{srg}, \mathbf{Z}_{\text{shape-skull}}) = \mathbf{S}_{\text{skull-plan}} - \mathbf{S}_{\text{skull-pre}}. \quad (6)$$

Critically, given that the \mathbf{Z}_{srg} simultaneously constrains the plausible transformation space for both the facial surface and skull during surgery, we combine it with the facial shape latent code $\mathbf{Z}_{\text{shape-face}}$ to deform the preoperative facial surface $\mathbf{F}_{\text{face-pre}}$ and predict the postoperative facial outcome:

$$\mathbf{F}_{\text{face-pred}} = \mathbf{F}_{\text{face-pre}} + \mathcal{F}_{srg}(\mathbf{x}_{\text{face}}, \mathbf{Z}_{srg}, \mathbf{Z}_{\text{shape-face}}). \quad (7)$$

The resulting $\mathbf{F}_{\text{face-pred}}$ is our framework’s prediction of the patient’s postoperative facial surface morphology.

3.5 Losses

Shape Module To supervise the learning of the shape latent space for facial surface (denoted as face), maxilla (mxl), and mandible (mdb), we design a composite loss function that enforces geometric fidelity, regularizes the latent space, and constrains anatomical landmarks. We jointly train the shape latent space by minimizing the overall shape loss as follows:

$$\mathcal{L}_{\text{shape}} = \lambda_{\text{glob}} \mathcal{L}_{\text{glob}} + \sum_{r \in \{\text{face, mxl, mdb}\}} \mathcal{L}_{\text{SDF}}^r + \lambda_{\text{lmk}}^r \mathcal{L}_{\text{lmk}}^r, \quad (8)$$

where $\mathcal{L}_{\text{glob}}$ is l_2 regularization term to constrain the norm of the global latent vector \mathbf{Z}_{glob} and $\mathcal{L}_{\text{lmk}}^r$ denotes the landmark regression loss for region $r \in \{\text{face, mxl, mdb}\}$, penalizing the deviation between predicted and ground truth landmarks. The region-specific SDF loss term $\mathcal{L}_{\text{SDF}}^r$ is defined as:

$$\begin{aligned} \mathcal{L}_{\text{SDF}}^r &= \lambda_{\text{surf}}^r |\mathcal{F}_{\text{shape}}(\mathbf{x}_r)| \\ &+ \lambda_{\text{surf}}^r \|\nabla \mathcal{F}_{\text{shape}}(\mathbf{x}_r) - \mathbf{n}(\mathbf{x}_r)\|_2 \\ &+ \lambda_{\text{eik}}^r [|\|\nabla \mathcal{F}_{\text{shape}}(\mathbf{x}_r)\|_2 - 1|] \\ &+ \lambda_{\text{far}} \exp(-\alpha \cdot |\mathcal{F}_{\text{shape}}(\mathbf{x}_r)|). \end{aligned} \quad (9)$$

The first term encourages the predicted SDF values at or near the surface to be close to zero. The second term enforces consistency between the gradients of the predicted SDF and the ground truth surface normals $\mathbf{n}(\mathbf{x}_r)$. The third term imposes the Eikonal constraint by regularizing the gradient magnitude to be close to 1, which is essential for SDF correctness. The final term penalizes near-zero SDF predictions at far-field locations by applying an exponential penalty, thereby promoting well-separated SDF values away from the surface.




Figure 4: Illustration of facial sub-region division for quantitative evaluation.

Surgery Module To effectively supervise the learning of a plausible deformation latent space that captures the complex biomechanical changes during orthognathic surgery, we design a loss function comprising correspondence consistency and deformation regularization for each anatomical region. We jointly optimize the surgery deformation latent space by minimizing the following composite loss:

$$\mathcal{L}_{surgery} = \lambda_{srg} \mathcal{L}_{srg} + \sum_{r \in \{\text{face, mxl, mdb}\}} \|\mathcal{F}_{srg}(\mathbf{x}_r) - \Delta(\mathbf{x}_r)\|_2. \quad (10)$$

where \mathcal{L}_{srg} is the latent regularization term that penalizes the l_2 norm of the surgery latent code \mathbf{Z}_{srg} to avoid overfitting and ensure compactness of the latent space. $\Delta(\mathbf{x}_r)$ is defined as: $\Delta(\mathbf{x}_r) = \mathbf{P}_{\text{gt-post}}^r - \mathbf{P}_{\text{gt-pre}}^r$, which represents the ground truth points displacement from preoperative to postoperative states for points \mathbf{x}_r .

4 Experiments and Results

4.1 Overview

In this section, we conduct comprehensive experiments to evaluate the performance of our neural implicit craniofacial model. Specifically, we separately evaluate: (1) the shape module for anatomical geometry reconstruction, (2) the surgery module for biomechanical deformation modeling, and (3) the postoperative facial appearance prediction.

4.2 Experimental Settings

Dataset We accessed archived CT scan data from the clinical archives of the Department of Oral & Craniomaxillofacial Surgery at a collaborating hospital. A total of 76 pairs of pre and post-surgery head CT images, along with multi-view facial scans, were collected, with 60 pairs used for training, 8 for testing and 8 for validation.

Data Preprocessing To standardize coordinate systems across pre- and post-surgery CT scans and ensure reliable spatial correspondences for our surgery module, we performed a data registration based on predefined template

face and skull models, following the approach in SCULPTOR (Qiu et al. 2022). The process includes an initial rigid alignment, followed by non-rigid deformation using embedded techniques to optimize vertex correspondence and anatomical landmarks. This results in meshes with consistent point and face counts, preserving topological consistency and integrity. Additional details are provided in the supplementary materials.

Baselines and Metrics We adopt the parametric model SCULPTOR (Qiu et al. 2022)—which jointly represents facial surface and bony structures—as the benchmark for evaluating our shape and surgery modules. For assessment of postoperative facial appearance prediction, we conduct comparative experiments with recent state-of-the-art methods in craniofacial surgery prediction including FSC-Net (Ma et al. 2022), ACMT-Net (Fang et al. 2024), and BPMNR (Yang et al. 2025). We reproduce them and use the same implementation details as mentioned in their papers.

For comprehensive comparisons, we use Point-to-Plane Distance (P2PL) and Chamfer Distance (CD) (Liu et al. 2020), both measured in millimeters (mm). Additionally, we provide a detailed regional analysis of postoperative facial appearance prediction accuracy, with Fig. 4 showing the divisions: (A) forehead, (B) nose, (C) lips, (D) chin, (E) left cheek and (F) right cheek.

Implementation Details We use 39, 43, and 16 anatomical landmarks for the facial surface, maxilla, and mandible, respectively. The global latent code \mathbf{Z}_{glob} has a dimension of 64. The local latent codes \mathbf{Z}_{loc} for the facial surface, maxilla, and mandible are 32, 64, and 32-dimensional, respectively. The surgery latent code \mathbf{Z}_{srg} is 512-dimensional. The local small MLP in Eq. (2) follows the DeepSDF architecture (Park et al. 2019) with 4 hidden layers, each containing 200 hidden units. Due to space constraints, additional network architectures and parameter settings are provided in the released code and the supplementary materials. The shape module is trained for 30,000 epochs with initial learning rates of 5×10^{-4} for network parameters and 1×10^{-3} for shape latent codes. Both are halved every 3,000 epochs. The surgery module is trained for 9,000 epochs with the same initial learning rates, decayed every 600 epochs. Optimization is performed using the Adam optimizer (Kingma and Ba 2014). All training is conducted on a single NVIDIA RTX 4090 GPU (24 GB).

4.3 Evaluation of Shape and Surgery Modules

We employ the parametric model SCULPTOR as a benchmark to evaluate the capabilities of both our shape and surgery modules. Specifically, for shape module evaluation, we utilize SCULPTOR’s shape component (denoted as SCULPTOR-S) for anatomical reconstruction, whereas for surgery module evaluation, we employ its trait component (denoted as SCULPTOR-T) to model surgical deformations.

Table 1 summarizes the geometry reconstruction and deformation modeling accuracy of our shape and surgery modules, respectively. Compared to SCULPTOR-S and SCULPTOR-T, our method significantly reduces both P2PL and CD errors across all regions. Fig. 5 illustrates visual

Metric	Shape module			Surgery module				
	Method	Face	Maxilla	Mandible	Method	Face	Maxilla	Mandible
P2PL↓	SCULPTOR-S	1.69±0.19	1.48±0.26	1.72±0.49	SCULPTOR-T	1.14±0.16	1.12±0.12	1.09±0.18
	Ours(single decoder)	0.62±0.16	0.44±0.05	0.54±0.21	Ours(global MLP _ψ)	0.76±0.17	0.48±0.07	0.42±0.08
	Ours(multi decoders)	0.52±0.10	0.42±0.05	0.51±0.17	Ours(separate MLP _ψ)	0.73±0.17	0.47±0.04	0.44±0.09
CD↓	SCULPTOR-S	4.73±0.34	4.00±0.42	3.88±0.86	SCULPTOR-T	3.69±0.28	3.19±0.24	2.73±0.28
	Ours(single decoder)	3.17±0.20	1.99±0.08	1.71±0.25	Ours(global MLP _ψ)	2.27±0.44	1.74±0.15	1.29±0.20
	Ours(multi decoders)	3.02±0.14	1.93±0.07	1.67±0.19	Ours(separate MLP _ψ)	2.25±0.38	1.67±0.10	1.35±0.19

Table 1: Quantitative comparison of the shape and surgery modules, including ablation studies conducted against SCULPTOR-S and SCULPTOR-T. Evaluation is performed on the face, maxilla, and mandible, measured in millimeters (mm). Lower values indicate better performance. The best and second-best performances are shown in **bold** and underline, respectively.

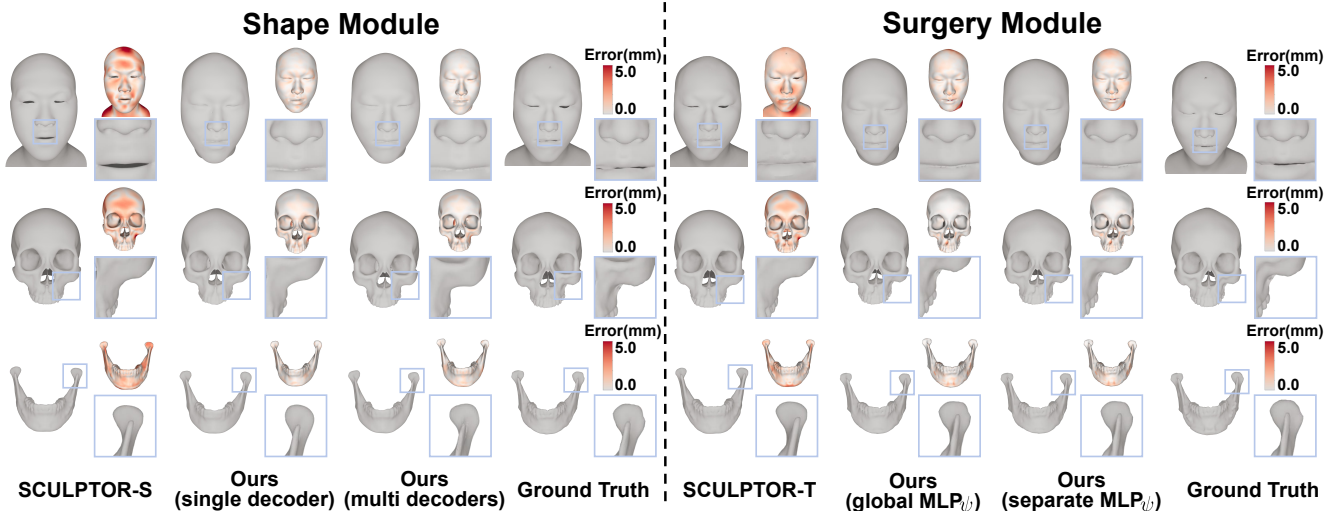


Figure 5: Qualitative performance comparison of the shape and surgery modules, including ablation studies conducted against SCULPTOR-S and SCULPTOR-T. Each row corresponds to a different anatomical region: facial surface (top), maxilla (middle), and mandible (bottom). In each column, the top-right subfigure shows the P2PL error map compared with ground truth (value range: 0-5 millimeters), while the bottom-right subfigure displays magnified views of specific anatomical structures.

comparisons of geometry reconstruction and surgery deformation against the ground truth. Our method produces smoother and more anatomically consistent surfaces, particularly in complex regions such as the lips, zygomatic bone, and temporomandibular joint (highlighted in the figure). The surgery module also better captures nonlinear deformations, yielding postoperative facial appearance predictions more closely aligned with the ground truth. Both quantitative and qualitative results demonstrate that our method significantly enhances anatomical reconstruction and surgical deformation modeling accuracy, laying a strong foundation for reliable postoperative facial appearance prediction.

4.4 Evaluation of Postoperative Facial Appearance Prediction

We evaluate the performance of our framework on postoperative facial appearance prediction against state-of-the-art methods. As shown in Table 2, our method consistently achieves the lowest P2PL and CD errors across most facial sub-regions, with marked improvements in critical areas such as the lips and chin. Our method yields the lowest

P2PL on the entire face (0.88 ± 0.08 mm), and outperforms BPMNR notably in the lips (1.05 ± 0.40 vs. 1.12 ± 0.42 mm) and chin (1.48 ± 0.36 vs. 1.62 ± 0.69 mm). It also attains the best CD overall (2.65 ± 0.12 mm), indicating robust surface-level accuracy across regions. These results underscore the advantage of our anatomically informed, region-specific deformation modeling. Fig. 6 showcases the qualitative P2PL error maps across representative cases, where our method exhibits notably fewer high-error regions and a more consistent error distribution, demonstrating superior accuracy in anatomically complex areas. Facial appearance prediction for a single patient on an RTX 4090 requires around 150 seconds in total—shape latent codes fitting (~ 70 s), surgery latent codes fitting (~ 50 s), and mesh generation (~ 30 s)—with a peak memory footprint of around 1.25 GB.

4.5 Ablation Studies

We further perform ablation studies on the internal architecture components of both the shape and surgery module, to assess their contributions to the overall performance.

For the shape module, we compare two configurations:

Metric	Method	A. Forehead	B. Nose	C. Lips	D. Chin	E. Cheek (L)	F. Cheek (R)	Entire face
P2PL ↓	FSC-Net	2.58±1.19	1.90±0.62	2.52±0.83	2.74±0.92	2.55±0.66	2.77±0.81	2.35±0.71
	ACMT-Net	1.12±0.34	1.32±0.32	1.56±0.49	2.31±0.83	1.21±0.31	1.48±0.32	1.29±0.26
	BPMNR	0.74±0.23	<u>0.97±0.36</u>	<u>1.12±0.42</u>	1.62±0.69	1.07±0.24	1.03±0.22	<u>0.93±0.14</u>
	Ours (global MLP _ψ)	<u>0.77±0.27</u>	<u>0.87±0.33</u>	<u>1.67±0.82</u>	<u>1.60±0.53</u>	1.17±0.37	1.41±0.44	1.03±0.21
	Ours (separate MLP _ψ)	<u>0.78±0.16</u>	0.75±0.17	1.05±0.40	1.48±0.36	<u>1.15±0.23</u>	<u>1.23±0.44</u>	0.88±0.08
CD ↓	FSC-Net	5.51±2.05	4.37±1.02	5.04±1.31	6.44±1.52	6.15±1.07	6.61±1.18	5.23±1.18
	ACMT-Net	3.18±0.65	3.37±0.56	3.84±1.04	5.83±1.30	4.21±0.51	4.35±0.88	3.58±0.51
	BPMNR	2.50±0.39	2.88±0.58	<u>2.89±0.66</u>	4.40±1.27	3.70±0.38	3.57±0.71	2.93±0.21
	Ours (global MLP _ψ)	<u>2.42±0.45</u>	<u>2.38±0.62</u>	<u>3.76±1.53</u>	4.33±0.90	<u>3.62±0.67</u>	3.98±0.86	<u>2.90±0.40</u>
	Ours (separate MLP _ψ)	2.38±0.25	2.27±0.24	2.61±0.72	4.18±0.62	3.61±0.33	<u>3.76±0.69</u>	2.65±0.12

Table 2: Quantitative results compared with FSC-Net, ACMT-Net, and BPMNR, including ablation studies on postoperative facial appearance prediction accuracy, measured in millimeters (mm). Lower values indicate better performance. The best and second-best performances are shown in **bold** and underline, respectively.

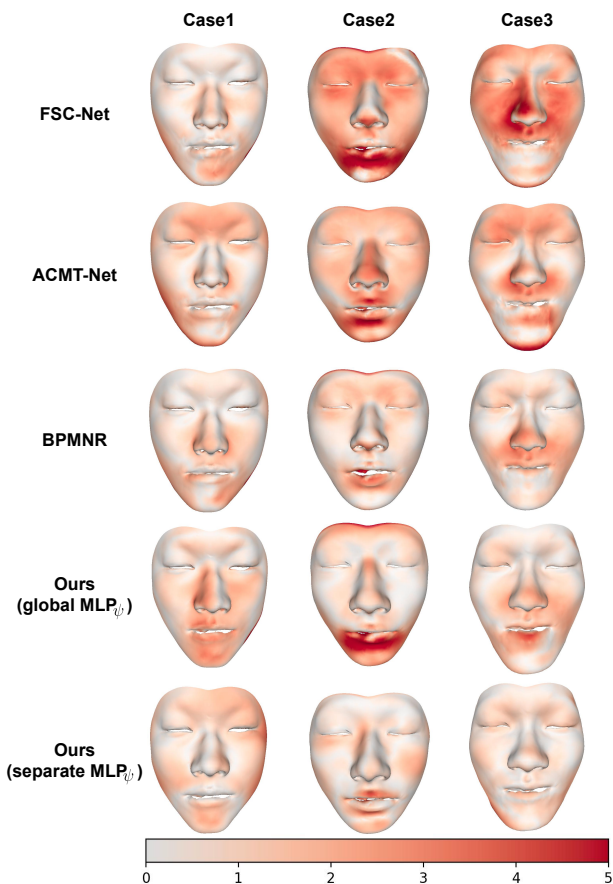


Figure 6: Qualitative P2PL error maps in comparison with FSC-Net, ACMT-Net, and BPMNR, including ablation studies on postoperative facial appearance prediction accuracy (value range: 0-5 millimeters).

one employing three region-specific SDF decoders (denoted as multi decoders), and the other using a single shared decoder (single decoder). This evaluates whether anatomical

specialization improves geometric reconstruction. For the surgery module, we assess the role of the dimensionality reduction network MLP_ψ. Specifically, we compare a shared global MLP_ψ that processes all region-specific latent codes jointly with separate MLP_ψ networks, each handling its corresponding region. This investigates whether region-specific modeling enhances deformation prediction.

The ablation results in Table 1, and Table 2 show that region-specific designs consistently improve the performance of both shape reconstruction and surgery deformation modeling. For the shape module, using separate SDF decoders better captures anatomical variations. In the surgery module, adopting region-specific MLP_ψ leads to more accurate deformation modeling across facial components. As shown in Fig. 5 and Fig. 6, these configurations produce smoother surfaces and more precise geometry, particularly in complex regions such as the lips, chin, and temporomandibular joint. These results highlight the benefit of anatomical decomposition in enhancing model fidelity.

5 Conclusion

In this work, we presented Neural Implicit Craniofacial Model (NICE) for accurate postoperative facial appearance prediction in orthognathic surgery. By leveraging region-specific implicit SDF decoders and shape latent representations, our shape module enables high-fidelity reconstruction of both facial soft tissue and bony structures. The surgery module further introduces anatomically guided, region-specific deformation decoders driven by a unified surgery latent code, effectively capturing complex nonlinear biomechanical responses of facial tissues to skeletal movements. Comprehensive experiments on clinical data demonstrate that our method outperforms existing state-of-the-art approaches in both geometric accuracy and anatomical consistency, particularly in critical regions such as the lips and chin. Future work will enhance latent-space interpretability and control via landmark-guided constraints, aiming to enable more structured and anatomically meaningful surgical outcome control.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant no. 62571328. The experiments of this work were supported by the core facility Platform of Electronics, SIST, ShanghaiTech University.

References

- Bao, J.; Zhang, X.; Xiang, S.; Liu, H.; Cheng, M.; Yang, Y.; Huang, X.; Xiang, W.; Cui, W.; Lai, H.; et al. 2024. Deep Learning-Based Facial and Skeletal Transformations for Surgical Planning. *Journal of Dental Research*, 103(8): 809–819.
- Fang, X.; Kim, D.; Xu, X.; Kuang, T.; Lampen, N.; Lee, J.; Deng, H. H.; Liebschner, M. A.; Xia, J. J.; Gateno, J.; et al. 2024. Correspondence attention for facial appearance simulation. *Medical Image Analysis*, 93: 103094.
- Genova, K.; Cole, F.; Sud, A.; Sarna, A.; and Funkhouser, T. 2020. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4857–4866.
- Giebenhain, S.; Kirschstein, T.; Georgopoulos, M.; R^unz, M.; Agapito, L.; and Nießner, M. 2023. Learning neural parametric head models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21003–21012.
- Huang, X.; He, D.; Li, Z.; Zhang, X.; and Wang, X. 2025. Maxillofacial bone movements-aware dual graph convolution approach for postoperative facial appearance prediction. *Medical Image Analysis*, 99: 103350.
- Kim, D.; Kuang, T.; Rodrigues, Y. L.; Gateno, J.; Shen, S. G.; Wang, X.; Deng, H.; Yuan, P.; Alfi, D. M.; Liebschner, M. A.; et al. 2019. A new approach of predicting facial changes following orthognathic surgery using realistic lip sliding effect. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22*, 336–344. Springer.
- Kim, D.; Kuang, T.; Rodrigues, Y. L.; Gateno, J.; Shen, S. G.; Wang, X.; Stein, K.; Deng, H. H.; Liebschner, M. A.; and Xia, J. J. 2021. A novel incremental simulation of facial changes following orthognathic surgery using FEM with realistic lip sliding effect. *Medical image analysis*, 72: 102095.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Knoops, P. G.; Borghi, A.; Ruggiero, F.; Badiali, G.; Bianchi, A.; Marchetti, C.; Rodriguez-Florez, N.; Breakey, R. W.; Jeelani, O.; Dunaway, D. J.; et al. 2018. A novel soft tissue prediction methodology for orthognathic surgery based on probabilistic finite element modelling. *PloS one*, 13(5): e0197209.
- Lampen, N.; Kim, D.; Xu, X.; Fang, X.; Lee, J.; Kuang, T.; Deng, H. H.; Liebschner, M. A.; Xia, J. J.; Gateno, J.; et al. 2023. Spatiotemporal incremental mechanics modeling of facial tissue change. In *International conference on medical image computing and computer-assisted intervention*, 566–575. Springer.
- Li, T.; Bolkart, T.; Black, M. J.; Li, H.; and Romero, J. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.*, 36(6): 194–1.
- Liu, M.; Sheng, L.; Yang, S.; Shao, J.; and Hu, S.-M. 2020. Morphing and sampling network for dense point cloud completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 11596–11603.
- Ma, L.; Lian, C.; Kim, D.; Xiao, D.; Wei, D.; Liu, Q.; Kuang, T.; Ghanbari, M.; Li, G.; Gateno, J.; et al. 2023. Bidirectional prediction of facial and bony shapes for orthognathic surgical planning. *Medical image analysis*, 83: 102644.
- Ma, L.; Xiao, D.; Kim, D.; Lian, C.; Kuang, T.; Liu, Q.; Deng, H.; Yang, E.; Liebschner, M. A.; Gateno, J.; et al. 2022. Simulation of postoperative facial appearances via geometric deep learning for efficient orthognathic surgical planning. *IEEE transactions on medical imaging*, 42(2): 336–345.
- Nadjmi, N.; Defrancq, E.; Mollemans, W.; Van Hemelen, G.; and Bergé, S. 2014. Quantitative validation of a computer-aided maxillofacial planning system, focusing on soft tissue deformations. *Annals of maxillofacial surgery*, 4(2): 171–175.
- Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 165–174.
- Posnick, J. C. 2021. Orthognathic surgery: Past–present–future. *Journal of Oral and Maxillofacial Surgery*, 79(10): 1996–1998.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qiu, Z.; Li, Y.; He, D.; Zhang, Q.; Zhang, L.; Zhang, Y.; Wang, J.; Xu, L.; Wang, X.; Zhang, Y.; et al. 2022. SCULPTOR: Skeleton-consistent face creation using a learned parametric generator. *ACM Transactions on Graphics (TOG)*, 41(6): 1–17.
- Stokbro, K.; Aagaard, E.; Torkov, P.; Bell, R.; and Thygesen, T. 2014. Virtual planning in orthognathic surgery. *International journal of oral and maxillofacial surgery*, 43(8): 957–965.
- Yang, J.; He, D.; Tian, X.; Ma, L.; Yu, J.; Wang, X.; and Zhang, Y. 2025. Blending Parametric Model and Neural Refinement: A Coarse-to-Fine Approach for Predicting Facial Changes in Orthognathic Surgery. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Zheng, Z.; Huang, H.; Yu, T.; Zhang, H.; Guo, Y.; and Liu, Y. 2022. Structured local radiance fields for human avatar modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15893–15903.