FracFormer: Fracture Reduction Planning with Transformer-Based Shape Restoration and Fracture Data Simulation

Sutuke Yibulayimu, Yanzhen Liu, Yudi Sang, *Member, IEEE*, Jingjiang Qin, Chao Shi, Chendi Liang, Gang Zhu, Yu Wang, Chunpeng Zhao, and Xinbao Wu

***Abstract*— Accurate orthopedic fracture reduction planning is essential for ensuring successful postoperative recovery and improving patient outcomes. However, current automatic methods are challenged by the complex and irregular fracture geometries and the scarcity of annotated training data. To address these challenges, we propose a novel approach that integrates learning-based shape restoration and fracture simulation. A transformer-based model is developed, which utilizes patch-to-patch shape translation and recursive fragment registration to iteratively refine fracture reduction poses. A deformable fracture generation model (DFGM) combines statistical shape modeling with clinically representative fracture patterns to generate diverse and realistic datasets, reducing the dependence on annotated samples. Tested on extensive clinical data with hipbone, sacrum, and femoral shaft fractures, the proposed method achieved mean errors of 1.85 mm and 3.40°, outperforming both template-based and existing learning-based methods. In addition, models trained solely on DFGM-synthesized data presented strong generalizability to real clinical data. The ablation experiments demonstrate the effectiveness of the fragment-aware network pipeline and the synthesis steps. Finally, a cadaver study with ground truth derived from the pre-injury scan further validated the performance of the method.**

***Index Terms*— Bone fracture, Surgical planning, Fracture simulation, Transformer, Point cloud deep learning**

I. Introduction

Fractures are among the most severe orthopedic injuries, particularly when involving multiple fragments or complex patterns, often leading to high disability and mortality rates [1], [2]. Surgical intervention is essential to restore anatomical structure, but conventional open reduction requires large incisions, resulting in significant blood loss and increased risks of nerve or vascular damage and infections [3]. Advances in intraoperative imaging and minimally invasive techniques with robot assistance have made image-guided closed reduction a more effective and less invasive alternative, offering reduced blood loss, lower infection rates, and faster recovery [4]. The success of closed reduction emphasizes the need for reliable and automated planning methods to determine the target pose of each bone fragment [5].

To improve surgical planning efficiency and accuracy, clinicians often employ 3D visualization and interactive manipulation tools [6], [7]. However, complex fractures—including impacted, severely displaced, or multifragmentary cases—can require hours of iterative adjustment [8], [9]. Furthermore, the final outcome depends heavily on the surgeon's expertise, complicating efforts to ensure consistently high-precision reductions [10]. In response, computational methods have been introduced. Geometry-based techniques use local shape features for fragment alignment but often fail when fracture surfaces are irregular or incomplete [11]–[14]. Template-based strategies rely on contralateral mirroring or reference shapes, but may fail in cases with bilateral injuries or natural anatomical asymmetry [14]–[18]. Statistical shape models (SSMs) capture anatomical variability across patients [19]–[22], but joint optimization of fragment poses and shape parameters remains challenging.

Recent data-driven deep learning methods have shown preliminary success in automatic reduction planning by training on a combination of simulated and real fracture data [9]. However, significant discrepancies persist when comparing such data to real clinical fractures, which often stem from high-energy trauma and involve severe deformations and comminuted fragments that go beyond straightforward displacement, as shown in Fig. 1. In addition, the considerable variability in fragment number, shape, and pose can hinder effective feature extraction, reducing the robustness of standard convolutional networks in complex scenarios.

Motivated by these challenges in effective learning of fracture morphologies and statistical modeling of fracture data, we present a novel approach that combines learning-based shape

This work was supported by the National Key Research and Development Program of China (Grant 2022YFC2504304), the Beijing Science and Technology Project (Grant Z221100003522007), and the National Natural Science Foundation of China. (Grant NSFC6247010104) *(Corresponding authors: Yudi Sang and Yu Wang.)*

Sutuke Yibulayimu, Yanzhen Liu, Jingjiang Qin, Chao Shi, Chendi Liang, and Yu Wang are with the Key Laboratory of Biomechanics and Mechanobiology (Beihang University), Ministry of Education, Beijing Advanced Innovation Center for Biomedical Engineering, School of Biological Science and Medical Engineering, Beihang University, Beijing, China (e-mail: sutuk@buaa.edu.cn; yanzhenliu@buaa.edu.cn; 21374024@buaa.edu.cn; 13361435056@163.com; liangchendi1997@sina.com; wangyu@buaa.edu.cn).

Chunpeng Zhao and Xinbao Wu are with the Department of Orthopaedics and Traumatology, Beijing Jishuitan Hospital, Capital Medical University, Beijing, China (e-mail: zcpllza2008@163.com; wuxinbao_jst@126.com).

Yudi Sang and Gang Zhu are with Beijing Rossum Robot Technology Co., Ltd., Beijing, China (e-mail: sangyudi@rossumrobot.cn; zhugang@rossumrobot.cn).

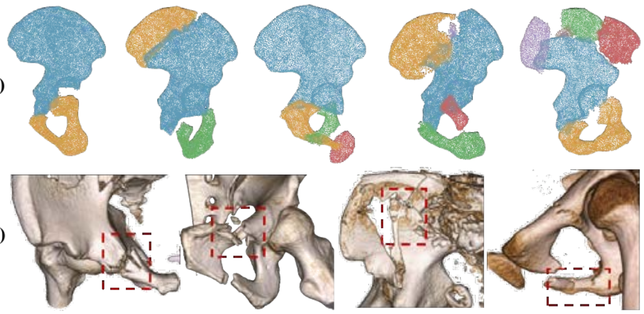

Fig. 1. Challenges in modeling and learning fracture features. (a) Variability in bone shape, fragment count and configurations. (b) Complex fracture surface presentations.

restoration with fracture simulation. We develop FracFormer, a transformer-based point cloud network to restore intact bone shapes from fractured data using fragment-aware patch encoding and patch-to-patch translation. A sparse-to-dense recursive refinement strategy further enhances alignment accuracy, effectively handling diverse fracture configurations. To support network training, we introduce a deformable fracture generation model (DFGM) that synthesizes diverse and realistic fractures based on statistical shape modeling, minimizing the reliance on large annotated datasets. Extensive experiments on clinical fracture cases demonstrate the high accuracy and robustness of our approach, proofing a reliable and generalizable solution for automated fracture reduction planning. Our source code is available at https://github.com/Sutuk/FracFormer.

II. Related Work

A. Fracture Reduction Planning

Existing reduction planning approaches can be broadly categorized into fracture surface matching and template-based methods. Surface matching aligns fragments based on geometric features but often fails when fracture surfaces are inaccurately segmented or defined [11]–[14]. Mirror-based approaches [14]–[17], as a template-based method, are useful for unilateral fractures but cannot handle bilateral fractures or bones lacking contralateral references. They also risk misalignment due to anatomical asymmetries [18]. SSMs have been explored as an alternative, enabling adaptive deformations of template to capture anatomical variability [19]–[22]. However, jointly optimizing fragment poses and deformation parameters often leads to suboptimal solutions in complex fractures.

Deep learning offers new possibilities for reduction planning but remains limited in practice. For instance, a PointNet-based framework has been proposed to establish correspondences between fracture fragments and deformable templates [23], but is limited to initial matching to the template and does not address the fitting or optimization process. A convolutional pose estimation network addressed sacroiliac joint dislocation [24] but cannot handle the more complex single-bone fractures. Recently, methods learning non-rigid transformations between healthy and fractured morphologies [9] have been proposed but struggled with large displacements due to the limitations of standard convolutional operations, underscoring the need for tailored feature extraction techniques in highly variable fracture scenarios.

B. Point Cloud-Based Shape Analysis

Detailed understanding of 3D shapes is pivotal for fracture reduction, as accurate geometric modeling directly influences alignment outcomes. Point-based neural networks have gained popularity for their flexibility in capturing local details without the constraints of volumetric grids or meshes. PointNet [25] introduced a permutation-invariant architecture, while dynamic graph convolution neural network (DGCNN) [26] enhanced local feature extraction through edge convolution. More recent transformer-based methods employ multi-head attention to learn long-range dependencies [27], [28]. In shape completion tasks, approaches like PoinTr [29] divide shapes into patches for finer-grained analysis and effective reconstruction. In the medical domain, Point2SSM integrates DGCNN and attention to accommodate irregular and patient-specific morphologies [30]. These advancements form a good foundation for point cloud learning-based fracture reduction planning.

C. Fracture Simulation

Fracture simulation for brittle materials poses significant challenges, as it requires accurate modeling of stress concentration and crack propagation. Traditional high-fidelity methods, such as finite element [31], boundary element [32], and material point methods [33], often demand days or weeks per simulation, making them impractical for large-scale data generation. To reduce computational burden, procedural algorithms pre-compute fracture patterns using Voronoi diagrams [34], [35], perturbed level set functions [36], or pre-authored patterns [37], [38], but often yield overly regular fragments that deviate from reality. Recent studies incorporating physical constraints, such as elastic vibration modes [39], produce more realistic results, but still rely on simplified assumptions, limiting their clinical applicability.

In the orthopedic domain, a bone density-weighted random distance map approach has been proposed, combining anatomical structure and bone density information with watershed segmentation to simulate bone fracture [9]. However, this method relies on real fracture data for training, underscoring current limitations in realism and diversity. As shown in Fig. 1b, clinical fractures exhibit irregular, mismatched surfaces and comminuted fragments, reflecting both complex fracture patterns and shape variability. These challenges highlight the need for sophisticated simulation techniques that jointly generate anatomical shapes and fracture patterns to better capture real-world clinical scenarios.

III. Method

We aim to estimate the target pose of each bone fragment using point cloud data derived from fracture segmentation [40]. As shown in Fig. 2, a transformer-based point cloud processing network is developed to restore bone shape, and a DFGM is developed to synthesize fractured dataset for network training.

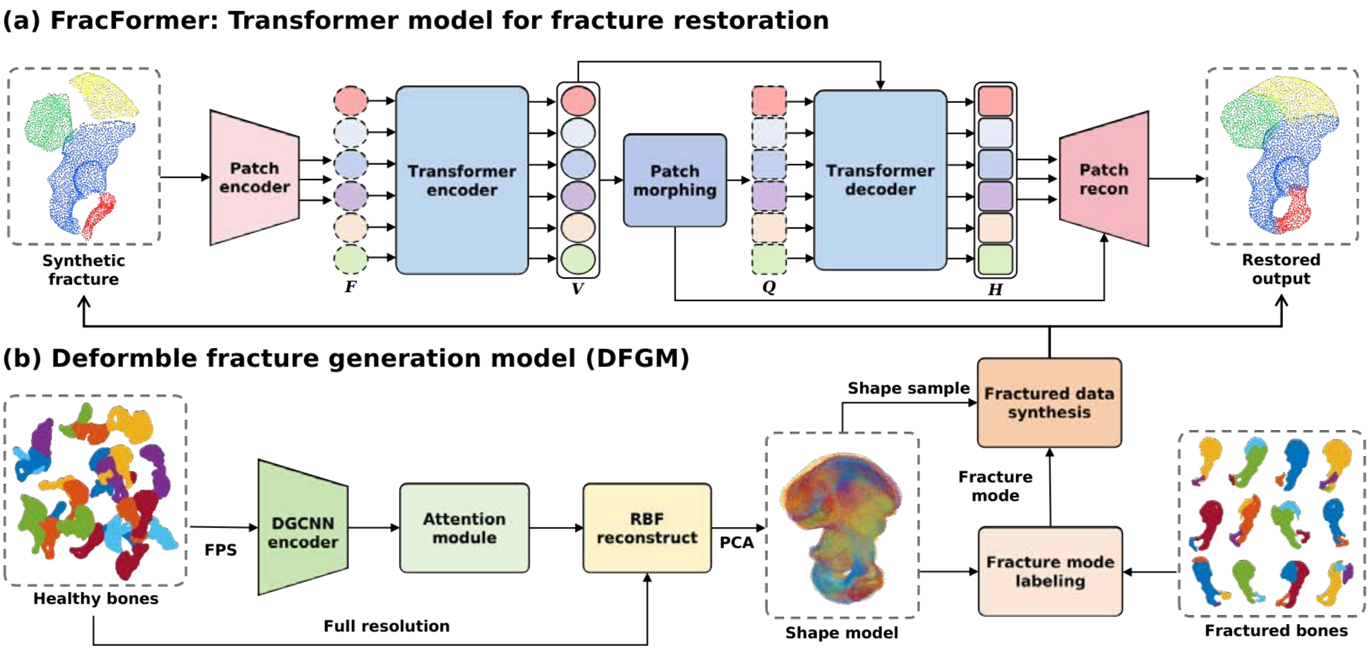

Fig. 2. Overview of the proposed method. (a) Bone shape restoration is performed by the FracFormer featuring fragment-aware patch encoding and patch morphing modules. (b) DFGM generates realistic fractured data for network training by modeling bone shapes and fracture modes.

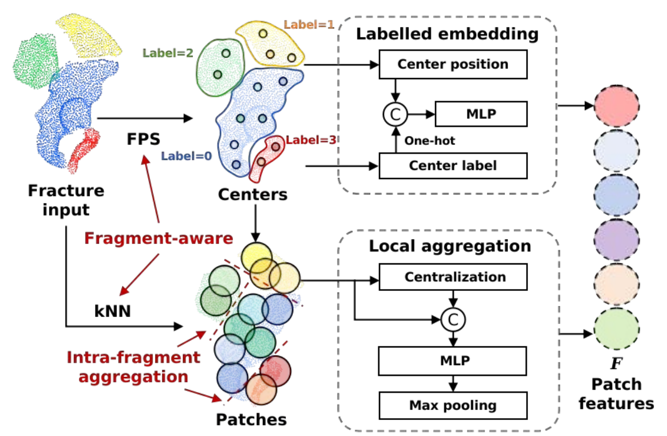

Fig. 3. The fragment-aware patch encoding module extracts local features with labeled embeddings for each fragment.

A. Transformer Model for Shape Restoration

Considering the variability in the number, pose, and interactions of fracture fragments, we formulate the fracture reduction planning task as a *patch-to-patch translation* problem based on transformer network, where the input is the labeled point cloud of fractured fragments, and the output is the point cloud of the reconstructed shape. A naive unified encoding for all fragments can lead to undesired interference among fragments. To mitigate this problem, we design FracFormer, a network pipeline that is aware of fragment labels throughout.

As shown in Fig. 2a, the proposed FracFormer first employs fragment-aware patch encoders to extract local geometric features while preserving fragment identity. These patch features are then processed by a transformer pipeline, capturing both intra- and inter-fragment relationships. A patch morphing module sits between the transformer encoder and decoder, refining patch centers and generating query embeddings. The transformer's decoder output, along with the refined centers, is fed into a final patch decoder to produce the fragment-wise restored shape. A tailored loss supervises the network training on two levels, ensuring accuracy and fragment-wise consistency in aligning sparse patch centers and reconstructing dense point clouds. Finally, the reconstructed patches guide a recursive alignment procedure, iteratively refining fragment poses until a stable anatomical reduction is achieved.

1) Fragment-Aware Patch Encoding: Fractured bones are represented by point clouds of multiple fragments. As shown in Fig. 3, to prepare the fragments for the transformer-based model, we first sample N_p patch centers across all fragments. Direct application of farthest point sampling (FPS) to the entire fractured bone may result in incomplete coverage near fragment boundaries. To address this, we allocate sampling centers proportionally to the point distribution of each fragment, and then apply FPS within each fragment to obtain center points and labels $\{c_i, \eta_i\}_{i=1}^{N_p}$ For each center c_i, we form a local patch P_i by selecting its k nearest neighbors (kNN) within the its entire fragment.

We applied a DGCNN to encode each patch P_i through dynamic graph construction and aggregation. Compared to the original model proposed in [26], both sampling and neighborhood aggregation are performed using the fragment-aware FPS and kNN methods described above. To embed fragment-specific identity, the patch center c_i is concatenated with its one-hot encoded fragment label η_i, and then processed by a multilayer perceptron (MLP):

$$F_i = \mathrm{DGCNN}(P_i) + \mathrm{MLP}([c_i, \text{one-hot}(\eta_i)]). \tag{1}$$

After processing all N_p patch centers in this manner, we obtain a unified set of patch features $\mathcal{F} = \{F_1, F_2, \ldots, F_{N_p}\}$, which is used as input to the subsequent transformer modules.

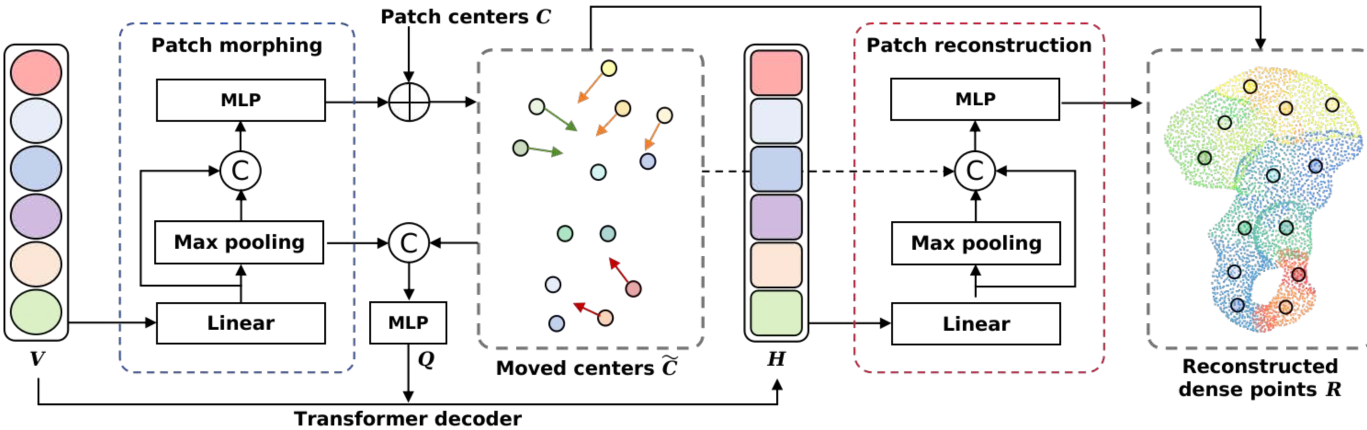

Fig. 4. The patch morphing module generates query embeddings and deformations of the sparse patch centers. The patch reconstruction module estimates the dense reconstructed point set from the moved centers and the decoder output.

2) Patch-to-Patch Translation with Transformer: In this formulated patch-to-patch translation task, the input fractured patch features are transformed into reconstructed patch features. This process is defined as:

$$\mathcal{V} = \mathrm{ME}(\mathcal{F}), \quad \mathcal{H} = \mathrm{MD}(\mathcal{Q}, \mathcal{V}), \tag{2}$$

where ME and MD are the encoder and decoder models, $\mathcal{V} = \{V_1, V_2, \ldots, V_{N_p}\}$ are the encoded patch features, $\mathcal{Q} = \{Q_1, Q_2, \ldots, Q_{N_p}\}$ are the query embeddings, and $\mathcal{H} = \{H_1, H_2, \ldots, H_{N_p}\}$ are the reconstructed patch features.

The architecture is based on the geometry-aware transformer proposed in [29]. The encoder employs six multi-head self-attention layers to model local and global dependencies among input patches. The decoder, composed of eight self-attention and cross-attention layers, aligns the query embeddings with the encoded features to reconstruct the fracture geometry.

3) Sparse-to-Dense Shape Reconstruction: Conventional encoder–decoder networks often lose fine-grained details in their global latent codes, and the reconstructed patches fail to maintain correspondence with the input patches. We address these problems by introducing more explicit guidance into the reconstruction process. Specifically, a patch morphing module is incorporated between the encoder and the decoder to predict both patch center deformations and query embeddings $\mathcal{Q}$ to guide subsequent reconstruction.

As shown in Fig. 4, the encoder output $\mathcal{V}$ undergoes linear projection followed by max-pooling to extract global context. The features are concatenated with patch features and passed through an MLP to estimate the deformations of patch centers $\mathcal{C}$. This yields the moved centers $\tilde{\mathcal{C}}$ and query embeddings $\mathcal{Q}$:

$$\tilde{\mathcal{C}} = \mathrm{Proj}\Big([\mathrm{Linear}(\mathcal{V}),\ \mathrm{Pool}\big(\mathrm{Linear}(\mathcal{V})\big)]\Big) + \mathcal{C}, \tag{3}$$

$$\mathcal{Q} = \mathrm{MLP}\Big([\tilde{\mathcal{C}},\ \mathrm{Pool}(\mathrm{Linear}(\mathcal{V}))]\Big). \tag{4}$$

The refined centers $\tilde{\mathcal{C}}$ serve as local anchors for reconstructing dense point cloud, while the output of the transformer decoder $\mathcal{H}$ provides local geometric details. For each refined center $\tilde{c}_i$, we infer the densely reconstructed point set from:

$$R_i = \mathrm{MLP}\Big([\mathrm{Linear}(H_i), \mathrm{Pool}(\mathrm{Linear}(\mathcal{H})), \tilde{c}_i]\Big), \tag{5}$$

where R_i denotes the reconstructed dense points around $\tilde{c}_i$. Since the patch centers $\mathcal{C}$ maintain correspondence with the original fragment labels, the reconstructed patches $\mathcal{R}$ preserve fragment-level consistency.

4) Loss Functions: The reconstruction process is supervised on sparse and dense levels. On the sparse level, mean squared error (MSE) aligns predicted patch centers with their corresponding ground-truth locations. While on the dense level, where point-wise correspondence is not established, Chamfer distance (CD) is used to enforce fine-grained consistency between the reconstructed point cloud and the dense ground truth.

For moved patch centers $\tilde{\mathcal{C}}$, reconstructed points $\mathcal{R}$, and the ground truths $\mathcal{C}^*$ and $\mathcal{R}^*$, the loss terms are:

$$\mathcal{L}_{\text{sparse}} = \frac{1}{N_p} \sum_{i=1}^{N_p} \|\tilde{c}_i - c_i^*\|_2^2, \tag{6}$$

$$\mathcal{L}_{\text{dense}} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \min_{r^* \in \mathcal{R}^*} \|r - r^*\|_2^2 + \frac{1}{|\mathcal{R}^*|} \sum_{r^* \in \mathcal{R}^*} \min_{r \in \mathcal{R}} \|r^* - r\|_2^2. \tag{7}$$

The total reconstruction loss is $\mathcal{L}_{\text{recon}} = \mathcal{L}_{\text{sparse}} + \mathcal{L}_{\text{dense}}$.

5) Recursive Inference and Registration: The reconstructed shape serves as a template for fracture reduction. During training, the network is exposed to fragments with various displacements that share the same target shape. During application, a recursive strategy iteratively refines fragment poses.

Specifically, at each iteration, coarse registration is achieved by aligning fragment centers $\mathcal{C}_i$ with the moved centers $\tilde{\mathcal{C}}_i$ via singular value decomposition (SVD). This is followed by fine registration, where each fragment is further registered to its densely reconstructed point set $\mathcal{R}_i$ via iterative closest point (ICP) algorithm. Finally, the transformations are applied, and the updated fragments serve as input to the next iteration.

This recursive process continues until the pose updates satisfy $\|T_i^{(\tau)} - I\| \leq \epsilon$ for all fragments or until reaching the maximum iteration $\tau_{\max}$. The final pose T_i^{total} is the cumulative product of transformations across all iterations. $T_i^{\text{total}} = \prod_{\tau=\tau_{\max}}^{1} T_i^{(\tau)}$.

B. Deformable Fracture Generation Model

As shown in Fig. 2b, to address the need for diverse training data for FracFormer, DFGM is proposed to learn anatomical variability and fracture patterns from a small set of fractured and healthy samples, and then generate diverse and clinically realistic fractures. It consists of three steps: First, anatomical shape modeling builds an SSM to capture anatomical variations and maintain consistent point correspondences, enabling fracture pattern propagation across anatomies. Second, fracture mode labeling maps clinical fracture patterns onto the SSM, creating a mode library that serves as the foundation for synthesizing new fracture cases. Finally, the fractured data synthesis step refines these modes by adjusting fragment boundaries, applying localized distortions, and introducing realistic fragment displacements, ensuring greater morphological diversity and clinical plausibility.

1) Anatomical Shape Modeling: To capture anatomical variability and consistently propagate fracture patterns across different bone shapes, an SSM is constructed for each bone type on healthy dataset. The process involves establishing point-wise correspondence among a set of bone geometries and applying principal component analysis (PCA) to model the underlying shape variability in a low-dimensional space.

Due to the computational burden of learning dense point correspondence directly, we start by processing a sparse subset of points. Given a set of N_s bone shapes $\mathcal{S} = \{S_1, S_2, \dots, S_{N_s}\}$, from each shape S_i, we first sample L sparse landmark points $X_i = \{x_i^\ell\}_{\ell=1}^{L} \in \mathbb{R}^{L\times 3}$ using farthest point sampling (FPS). To establish point-wise correspondence, these sparse points are processed by an unsupervised network inspired by Point2SSM++ [41]. A DGCNN-based encoder [26] extracts local and global features for each point, and an attention module computes an attention weight matrix $W_i \in \mathbb{R}^{M\times L}$, producing a rearranged points sequence $Y_i \in \mathbb{R}^{M\times 3}$ as a linear combination of input points in X_i:

$$Y_i = W_i \cdot X_i. \tag{8}$$

The network for predicting the attention map W is trained by minimizing the correspondence loss [41], defined as $\mathcal{L}_{\text{corr}} = \mathcal{L}_{\text{CD}} + \mathcal{L}_{\text{cons}} + \mathcal{L}_{\text{ME}}$, where the Chamfer distance loss $\mathcal{L}_{\text{CD}}$ ensures reconstruction fidelity, the consistency loss $\mathcal{L}_{\text{cons}}$ enforces pose invariance, and the pairwise mapping error $\mathcal{L}_{\text{ME}}$ regularizes correspondence across minibatch. As a result, the rearranged sparse point sequence Y_i is semantically consistent across shapes, with inherent correspondence.

Then, an arbitrary shape $S_r \in \mathcal{S}$, represented by its dense points is selected as reference to match all shapes in $\mathcal{S}$ through radial basis function (RBF) interpolation. Using the sparse points Y_r on S_r and Y_i on a candidate shape S_i, the rearranged dense point sequence is computed as:

$$\hat{S}_i = S_r + K(S_r, Y_r) \cdot K(Y_r, Y_r)^{-1} \cdot (Y_i - Y_r), \tag{9}$$

where $K(\cdot,\cdot)$ is the pairwise distance matrix between two point sets. Each element of K is the kernel distance between two points. A thin plate spline kernel $\varphi(d) = d^2 \log(d)$ is adopted, where d is the Euclidean distance.

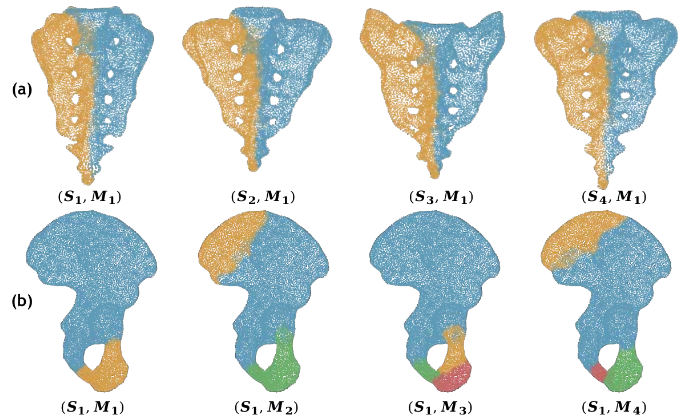

Fig. 5. Interplay between shape variation and fracture mode. (a) Sampling different shapes with the same fracture mode. (b) Sampling different fracture modes on the same shape.

This rearrangement enables a consistent parameterization of all shapes in $\mathcal{S}$. Each shape is then represented as stacked surface point coordinates. Generalized procrustes analysis (GPA) is performed to align these shapes with a common reference frame, and PCA is applied to model shape variability:

$$S = \bar{S} + \Phi b, \tag{10}$$

where $\bar{S}$ is the mean shape, Φ is the matrix of eigenvectors, and b is the vector of shape parameters. A wide range of anatomical variations can be simulated by adjusting b.

2) Fracture Mode Labeling: We aim to model the duel statistics of anatomical shapes and fracture patterns from real clinical data and synthesize fractures accordingly. However, since the distribution of the fracture surfaces is highly complex and discontinuous, we resort to mapping the fracture patterns from typical cases to the SSM and introduce rich variations during the fracture synthesis process. We establish a *fracture mode* representation on the SSM by labeling each point of the SSM based on fragment size rankings from fractured samples.

To align the SSM with a fractured bone, represented as N fragment shapes $\{A_n\}_{n=1}^{N}$, we aim to find the optimal shape parameters b and transformations $\mathcal{T} = \{T_n\}_{n=1}^{N}$. Adopting the iterative procedure described in [20], the optimization problem is formulated as:

$$(\mathcal{T}, \hat{b}) = \arg\min_{(\mathcal{T},b)} \left\| \bigcup_{n=1}^{N} T_n(A_n) - (\bar{S} + \Phi b) \right\|, \tag{11}$$

where eigenvalue constraints $|b_i| < 3\sqrt{\lambda_i}$ are applied to ensure that the deformed shape remains anatomically plausible [42].

Let $\{s_i\}_{i=1}^{p}$ be the set of p surface points on the deformed SSM corresponding to the estimated shape parameters $\hat{b}$. We define a fracture mode M as a discrete vector of length p:

$$M = (m_1, m_2, \dots, m_p)^T, \quad m_i \in \{1, \dots, N\}, \tag{12}$$

where each element m_i denotes the fragment index for the i-th point s_i. Formally,

$$m_i = \arg\min_{n\in\{1,\dots,N\}} \left\| s_i - T(A_n) \right\|, \tag{13}$$

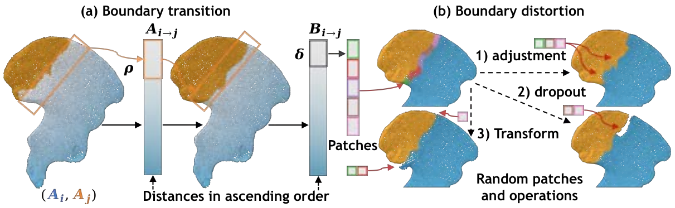

Fig. 6. Illustration of the boundary modification process. (a) Boundary transition through fragment label reassignment. (b) Boundary distortion through patch-based operations.

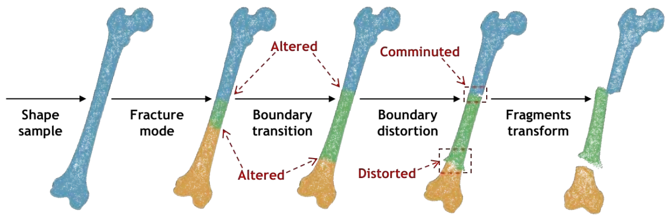

Fig. 7. Series of operations to synthesize realistic fractures.

which assigns each SSM point s_i to the nearest fragment $\mathcal{A}_n$. By labeling multiple fractured samples in this manner, a collection of modes is obtained.

As illustrated in Fig. 5a, a fracture mode M is randomly selected from the mode library and applied to different sampled SSM shapes $\{S_i\}$, ensuring that the synthesized fractures maintain the same clinically realistic fragment-partition pattern while adapting to varying anatomical structures. Similarly, as shown in Fig. 5b, different clinical fracture patterns can be represented within the same bone shape S, where applying various modes $\{M_i\}$ to a fixed shape results in distinct fragment partitions.

3) Fractured Data Synthesis: While fracture modes provide a representative set of patterns, they alone are insufficient to capture the rich morphological variability observed in clinical fractured data. Real fracture surfaces can exhibit missing regions (as comminuted bone shards are sometimes ignored in segmentation), irregular boundaries, and complex fragment interactions (such as collision and non-rigid deformation), that go beyond simple label partition. To address this, we introduce a series of operations, including transition, distortion, surface closure, and pose adjustments.

a) Boundary Transition: The transition process introduces variability in fragment boundaries, generating a broader range of fracture patterns. Each fracture mode is initially represented as a distribution of fragment labels on the SSM. The fragments A_i and A_j are defined by grouping points assigned to labels i and j, respectively. To modify fragment boundaries, we first identify adjacent fragment pairs based on spatial proximity. Given two such fragments, we calculate the distance of a point x in fragment A_i to fragment A_j as:

$$d(x, A_j) = \min_{x' \in A_j} \|x - x'\|. \tag{14}$$

Using this definition, the boundary region $B_{i\to j}$ between two adjacent fragments is identified as:

$$B_{i\to j} = \left\{x \in A_i \,\middle|\, d(x, A_j) \le \delta\right\}. \tag{15}$$

If $|B_{i\to j}|$ meets a predefined size threshold, we consider A_i and A_j to be *adjacent fragments*. Otherwise, the fragment pair is not included in the subsequent postprocessing operations.

As illustrated in Fig. 6a, to modify the boundary between adjacent fragments, we transfer a fraction of points from one fragment to the other. Taking $i \to j$ as an example, we first sort all points in A_i in ascending order of their distance $d(x, A_j)$, then reassign a fraction ρ of the nearest points. Formally, the reassigned set $A_{i\to j}$ is defined as:

$$A_{i\to j} = \left\{x \in A_i \,\middle|\, d(x, A_j) \le d_\rho\right\}, \tag{16}$$

$$d_\rho = \mathrm{Quantile}_\rho\big(\{d(x, A_j) \mid x \in A_i\}\big), \tag{17}$$

where d_ρ is the distance threshold corresponding to the top-ρ quantile of distances in A_i to A_j. For all points in $A_{i\to j}$, the fragment labels are updated to j. By iterating this process over all adjacent fragment pairs, we effectively slide fragment boundaries, altering the fracture mode mask while maintaining anatomical consistency.

b) Boundary Distortion: In addition to boundary transition, clinical fractures exhibit localized irregularities, missing shards, and subtle deformations. To capture these features, we subdivide each boundary region $B_{i\to j}$ into adaptive patches based on spatial distribution, as illustrated in Fig. 6b. Given a set of boundary points, we first identify the pair of farthest points in $B_{i\to j}$, denoted as (x_a, x_b). One of these points, x_a, is chosen as the reference anchor. Sorting all points in $B_{i\to j}$ in ascending order of their Euclidean distance from x_a, we partition them into K contiguous sub-patches $P_1, P_2, \dots, P_K$, where K is adaptively determined based on the total number of points in $B_{i\to j}$ while constrained within $2 \le K \le 6$ to ensure adequate patch size.

A subset of the patches $\{P_k\}$ is randomly selected, and each selected patch undergoes one of three possible operations, applied with a certain probability to mimic real fracture irregularities.

The first operation is patch-wise label reassignment, where all points in the selected patch P_k are transferred to the adjacent fragment, effectively altering the local boundary structure. This operation modifies the shape of the fragment while maintaining anatomical consistency.

The second operation applies a localized transformation to simulate deformation. For a randomly selected patch P_k, the deformation magnitude $\alpha(x)$ is determined as:

$$\alpha(x) = 1 - \frac{d(x, A_j) - d_{\min}}{d_{\max} - d_{\min}}, \tag{18}$$

where $d_{\min}$ and $d_{\max}$ are the minimum and maximum of $d(x, A_j)$ in P_k, ensuring stronger deformations near the adjacent fragment. A localized scale-shear transformation T_{distort} is then applied to all points in P_k:

$$x_{\text{distort}} = x + \alpha(x)\,\big(T_{\text{distort}}(x) - x\big), \quad \forall x \in P_k. \tag{19}$$

The third operation randomly removes all points within a selected patch P_k, simulating the loss of small bone fragments as seen in comminuted fractures.

These three operations are applied sequentially, with each step involving a new boundary patching process. For a given

TABLE I
SSM IMPLEMENTATION SETTINGS AND EVALUATION METRICS FOR THE THREE ANATOMIES.

Anatomy	Sample	Point	Fracture modes	PCA modes	Completeness(%)	Generality(mm)	Specificity(mm)
Hipbone	344×2	11,853	38	121	95.01	2.10 ± 0.18	4.15 ± 0.32
Sacrum	344	10,392	11	79	95.00	2.32 ± 0.63	3.80 ± 0.32
Femur	88×2	11,758	11	97	99.01	2.15 ± 0.29	3.84 ± 0.46

adjacent boundary, different operations are randomly applied, resulting in varied fracture modifications across iterations. This iterative approach progressively modifies the fracture morphology, producing irregular and discontinuous patterns that better reflect real clinical fractures.

c) Data Generation: As shown in Fig. 7, fractured samples are generated by uniformly sampling the SSM shape parameters b within $\pm 3\sqrt{\lambda_i}$ to create new anatomical geometries, followed by applying a random fracture mode with morphing and distortion to introduce boundary variability. Since the fragments are derived from intact surface models without predefined contact interfaces, a topology-based mesh refinement algorithm is applied to close the induced open surfaces and remain consistent with real clinical data [43].

The fragments are labeled according to the rank of their size. Each fragment, except the largest (label 0), is randomly assigned a clinically plausible transformation to generate the final fractured sample. The largest fragment remains stationary as a reference. The intact shapes without displacement are preserved as ground truth for network training.

IV. Experiments and Results

A. Experimental Setup

1) Data: We evaluated our method on three different bone anatomies: hipbone, sacrum, and femur. We selected 430 healthy hipbone and sacrum CT scans from the *CTPelvic1K* dataset [44], and 188 hipbone fracture and 52 sacrum fracture scans from the *PENGWIN* dataset [40]. In addition, we collected 110 healthy and 52 femoral shaft fracture scans from clinical data of Beijing Jishuitan Hospital from 2022 to 2024. The retrospective usage of the femoral CT was approved by the Ethics Committee of the institute (K2023-380-00).

The healthy samples were split into training and test sets in an 8:2 ratio to build the SSM of each anatomy. Due to the scarcity of fractured data, fracture samples were divided 2:8, with 20% used to construct fracture modes in DFPM and 80% reserved for evaluation. In total, 232 real fracture samples (150 hipbones, 41 sacra, and 41 femurs) were used for testing, making this one of the largest fracture test datasets to date. Fragment segmentation was performed using deep networks [40], followed by manual verification and adjustment.

Ground-truth reduction poses were annotated by two clinical experts, each with more than five years of surgical planning experience. To ensure consistency, the annotation process was performed on a CAD platform, following a standardized protocol, where all smaller fragments were aligned to the main fragment using translation and rotation operations. To assess variability among experts, two independent sets of annotations were generated by different experts, and inter-observer error was measured across the three bone anatomies to quantify annotation consistency.

2) Baseline Methods: We compared the proposed network model with several reduction planning methods, including average templates (Mean) [45], statistical shape templates (SSM) [20], mirrored templates (Mirror) [15], and fracture surface matching based on mirrored templates (Matching) [14]. Due to their reliance on the contralateral part, Mirror and Matching were restricted to unilateral hipbone and femur fractures only. For a fair and consistent comparison, those cases with contralateral reference were tested separately in our experiments. Furthermore, a transformer-based point cloud generation method, AdaPoinTr [46], was included as a learning-based baseline. To test the efficacy of the proposed fracture simulation method, models trained on DFPM-synthesized data were compared against models trained directly on real clinical fracture data and the breaking good method [39].

We also aimed to compare our method with the structural restoration network (SRN) proposed in [9]. However, due to challenges in re-implementing their approach and testing it within the same framework, we instead applied our method to the dataset used in their study and compared our results with those reported in their paper.

3) Evaluation: Reduction accuracy was assessed using translational error, defined as the Euclidean distance between the centers of reduced fragments and their ground truth, and rotational error, defined as the axis-angle difference between the planned and ground-truth pose matrices for each fragment. The shape after reduction was assessed using the L2 Chamfer distance (CD) relative to the ground truth. Additionally, part accuracy (PA) was calculated to measure the fragment-wise success rate of alignment. Following established definition [47], PA was determined by thresholding the CD between the normalized result and the normalized ground truth at 0.1.

4) FracFormer Implementation: The proposed method was implemented in PyTorch, on a platform equipped with an NVIDIA RTX 4070Ti GPU. The input point clouds of fractured bones were initially downsampled to 2,048 points using FPS. Fragment labels were assigned based on size ranking and encoded as one-hot vectors with seven dimensions, accommodating the maximum of six fragments observed in our datasets. The patch encoding module generated 256 patch tokens through four steps of 16 neighborhood point aggregations, each with 128-dimensional features. And the transformer used 384 dimensions. Reconstruction was performed at the sparse level with 256 points and the dense level with 4,096 points. AdamW optimizer was used with an initial learning rate of 10^{-4} and weight decay of 5×10^{-4}, following a cosine annealing schedule. The training spanned 300 epochs with a batch size of 20. The recursive process allowed up to 10

TABLE II
QUANTITATIVE PLANNING RESULTS ON HIPBONE AND FEMUR DATA THAT HAS CONTRALATERAL REFERENCE.

Method	Hipbone (76)				Femur (17)			
	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)
Mean [45]	6.07	5.25	4.55	46.05	5.33	12.64	8.45	62.75
SSM [20]	5.74	4.69	4.19	52.30	6.15	13.44	9.17	67.65
Mirror [15]	5.18	3.01	3.22	67.43	4.11	3.92	4.33	87.25
Matching [14]	6.40	3.08	3.46	67.32	4.33	4.84	5.40	86.27
AdaPoinTr [46]	5.08	3.04	3.13	69.74	7.59	5.91	5.65	82.35
FracFormer	**3.55**	**1.74**	**2.36**	**87.83**	**3.47**	**2.37**	**3.31**	**96.08**

TABLE III
QUANTITATIVE PLANNING RESULTS ON BILATERAL FRACTURE DATA WITHOUT CONTRALATERAL REFERENCE.

Method	Hipbone (74)				Sacrum (41)				Femur (24)			
	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)
Mean [45]	7.48	5.89	4.85	49.44	6.57	3.65	3.41	61.79	8.11	11.31	8.23	59.72
SSM [20]	6.74	5.29	4.50	49.21	6.19	3.60	3.39	61.79	6.68	10.65	7.89	57.64
AdaPoinTr [46]	6.24	3.16	2.95	74.44	5.74	2.44	2.58	81.71	7.68	5.96	5.64	75.00
FracFormer	**2.93**	**1.40**	**2.04**	**89.98**	**3.45**	**2.00**	**2.13**	**87.80**	**4.22**	**2.94**	**3.92**	**91.67**
Inter-observer	1.09	0.75	1.27	97.03	1.38	0.89	1.47	95.93	1.42	0.92	1.97	100.00

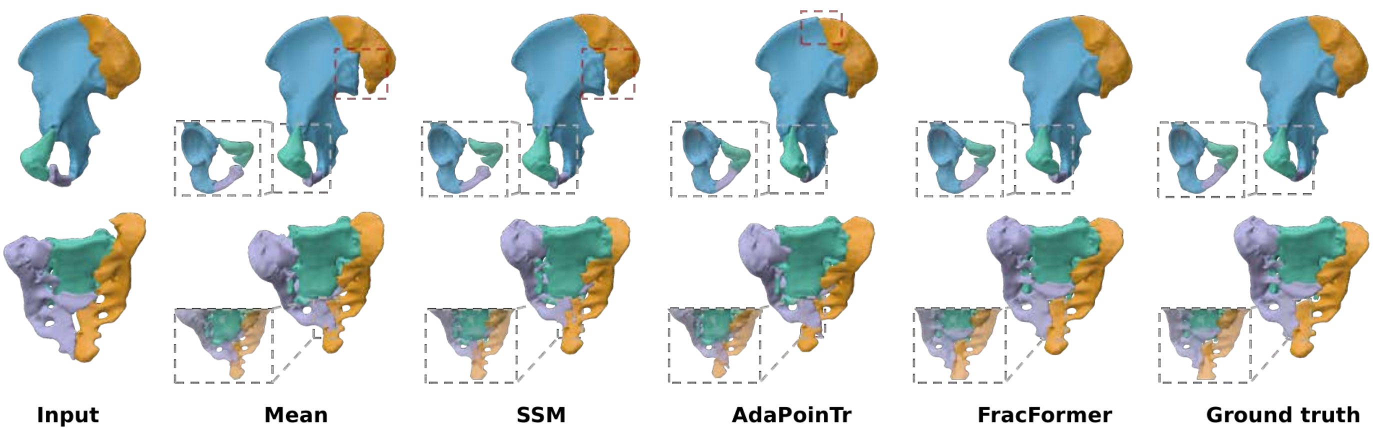

Fig. 8. Example reduction planning results on hipbone and sacrum.

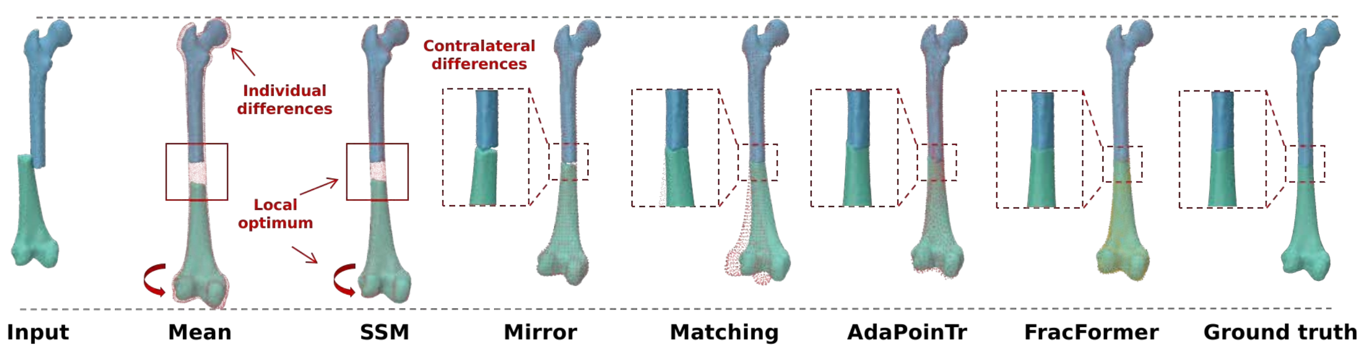

Fig. 9. Example femur reduction planning results in posterior view. The corresponding templates for reduction are shown in red. Alignment of the femoral linea aspera can be observed in the zoomed in view.

iterations and used a convergence threshold ϵ of 1.

5) DFGM Implementation: The SSM settings for hipbone, sacrum, and femur are summarized in Table I. The morphing step used a maximum distance threshold δ of 10 and randomly adjusted fragment proportions ρ between 1% and 30%. Boundary patch transformation and dropout occurred at a probability of 0.2. Distortion transformations were applied with independent scaling factors for each axis ranging from 0.6 to 1.4 and shear factors between ±0.4 for each axis. The synthesized fragments were further moved with random translation up to 15 mm and rotation up to 30° to simulate realistic displacement. For each anatomy, a total of 20,000

TABLE IV
QUANTITATIVE COMPARISON OF DFPM-SIMULATED DATA AGAINST CLINICAL CASES AND BREAKING GOOD IN TRAINING THE NETWORK.

Data source	Hipbone (150)				Sacrum (41)				Femur (41)			
	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)
Clinical cases	9.71	6.29	4.44	56.50	7.41	6.89	4.86	50.41	15.24	18.43	17.98	49.19
Breaking good [39]	9.42	5.94	3.97	58.61	6.89	4.21	3.48	60.16	9.57	11.42	9.51	60.98
DFGM	**3.24**	**1.57**	**2.20**	**88.89**	**3.45**	**2.00**	**2.14**	**87.80**	**3.91**	**2.70**	**3.67**	**94.31**

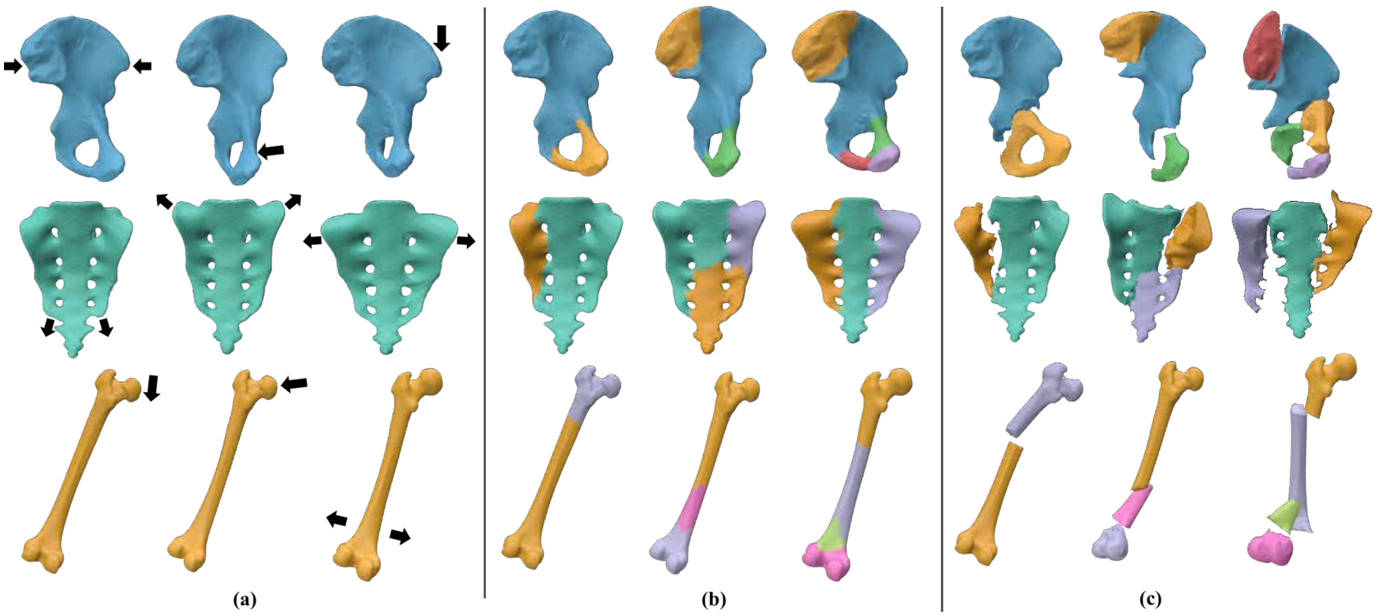

Fig. 10. Example simulated data from the DFGM. (a) Anatomical variations generated by the SSM along the first three principal components. (b) Different fracture modes applied on the generated anatomies. (c) Final simulated fractures with transition, distortion, and transformation applied.

training samples were synthesized before random transformations were applied. Properties of the SSM and synthetic dataset are analyzed in V-A and V-B.

B. Comparison of Reduction Planning Methods

Performance was evaluated on the group with contralateral reference and the group without, separately. As shown in Tables II and III, our method consistently outperformed all baselines. Paired t-tests confirmed that the improvements were statistically significant ($p < 0.05$) across all metrics and fracture types. Some example reduction planning results are shown in Fig. 8, and a detailed comparison of the femur results is further illustrated in Fig. 9. Mean and SSM struggled to capture individual variability. Mirror performed competitively by leveraging healthy templates, but was sometimes affected by inherent anatomical asymmetry, such as slight differences in femur anteversion or length. Matching often failed to accurately identify and align the fracture surfaces and thus lacked robustness. Despite being trained on the same datasets, AdaPoinTr resulted in less accurate shapes, possibly due to the lack of fragment-level consistency in its network pipeline.

The inter-observer variability remained lower than the prediction error of our method, confirming that the ground-truth annotations provide a stable reference for evaluating reduction accuracy.

Furthermore, we compared our method with the SRN method proposed by Zeng et al. [9], using the FP103 fracture subset of the *CTPelvic1K* dataset [44] following the same evaluation framework in the paper. Our method achieved lower rotational error of (2.27 ± 2.08)°, compared to (3.18 ± 2.59)° of SRN, and lower translational error of (2.03 ± 1.25) mm, compared to (2.88 ± 2.49) mm of SRN, demonstrating superior accuracy in reduction planning. Unlike our approach, which explicitly models fracture fragments represented by point clouds, SRN operates as a voxel-based network without fragment-level constraints, generating the overall reduced shape similarly to AdaPoinTr [46]. However, working in the voxel space directly may introduce additional computational complexity and limit fine-grained fragment alignment, potentially affecting its accuracy.

Our method achieved an average runtime of 2.03 seconds, a significant improvement over 42 minutes in the SSM-based method [19] and 22 seconds in the SRN method [9].

C. Comparison of Fracture Synthesis Methods

We evaluated the effectiveness of DFPM by comparing models trained with real clinical cases and the breaking good method [39]. As shown in Table IV, networks trained on simulated data outperformed those trained on real clinical fractures. Real datasets, limited in size and diversity, hindered the generalization of deep learning models, whereas the DFPM

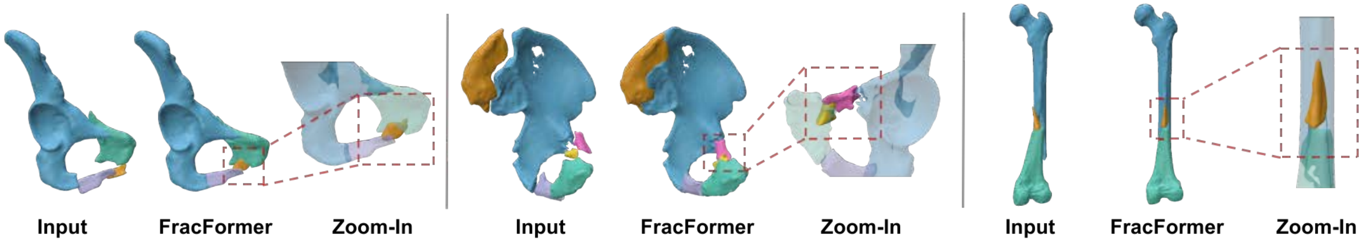

Fig. 11. Examples of repositioning error of small fragments.

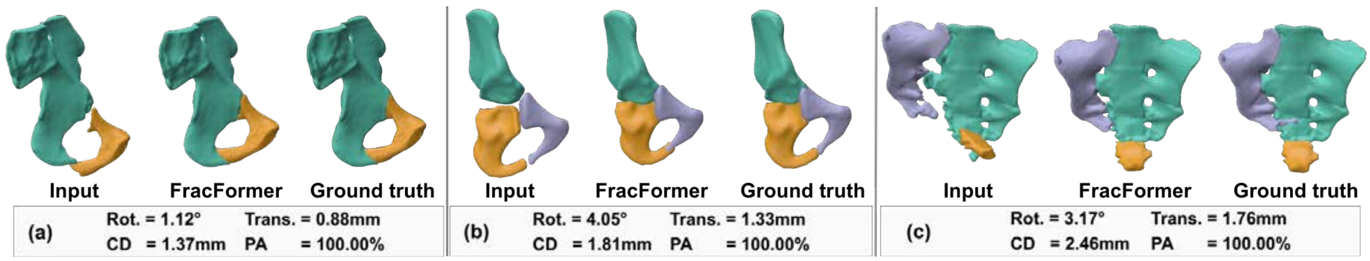

Fig. 12. Reduction planning results on three special cases. (a) Incomplete CT scan. (b) Pediatric fracture. (c) Fracture with severe deformity.

TABLE V

QUANTITATIVE PLANNING RESULTS IN THE ABLATIVE STUDY. THE UPPER SECTION SHOWS ABLATIONS ON THE FRACFORMER MODULES. THE LOWER SECTION SHOWS ABLATIONS ON THE DFGM.

Method	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)
Replace FAPE	6.61	3.30	3.25	68.44
Without label embed.	4.48	2.28	2.61	81.89
Remove FAPM	4.51	2.43	2.58	81.00
Without recursion	5.30	2.54	2.69	80.94
Without closure	6.39	5.28	4.01	52.56
Without transition	4.73	2.33	2.51	82.61
Without distortion	4.43	2.19	2.44	83.17
Full model	**3.24**	**1.57**	**2.20**	**88.89**

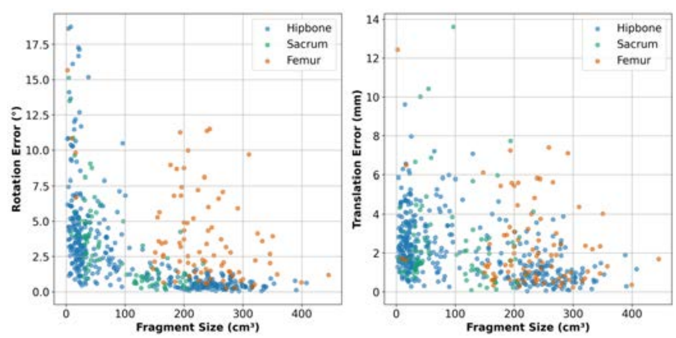

Fig. 13. Impact of fragment size on rotation and translation errors.

bridged this gap by generating abundant and realistic synthetic samples. In contrast, breaking good often produced fracture patterns that are clinically unrealistic. This underscores DFPM's advantage in capturing authentic fracture modes from minimal clinical data. Examples in Fig. 10 highlight DFPM's capability to generate diverse fracture patterns that align with clinical observations.

D. Ablation Study

The contributions of the network and DFPM components were analyzed in an ablative study on the hipbone dataset. As shown in Table V, replacing the fragment-aware patch encoding (FAPE) module with standard unified encoding caused significant performance degradation, highlighting its importance for fragment shape and pose differentiation. Incorporating fracture label embedding expedited convergence and ensured reconstruction consistency. The fragment-aware patch morphing (FAPM) module notably enhanced the network's ability to integrate fragment shapes and improved local reconstruction accuracy. Furthermore, removing the recursive strategy significantly reduced the alignment accuracy, underscoring its role in refining fragment poses.

In DFPM, the removal of key techniques such as fracture surface closure, boundary transition, and distortion led to noticeable drops in simulation quality. The removal of the closure of the fracture surface caused a significant gap between training data and the segmentation results from real clinical data. Removing fragment transition or distortion reduced data diversity, thus weakening model robustness and widening the generalization gap.

E. Impact of Fragment Size

To evaluate the impact of fragment size on reduction accuracy, we analyzed its relationship with both translation and rotation errors. As shown in Fig. 13, both errors decrease as fragment size increases, indicating that larger fragments contribute to more stable and precise reductions. Notably, rotation error is particularly sensitive to fragment size, showing a sharper increase for smaller fragments compared to translation error. These findings are consistent with previous studies on the error distribution of small fragments [9], [23].

Fig. 11 presents typical cases where the algorithm struggles to accurately reposition small fragments. The primary challenges came from segmentation inconsistencies and insufficient geometric features, which made precise alignment

TABLE VI
QUANTITATIVE REDUCTION PLANNING RESULTS IN THE CADAVER STUDY. THE BEST PERFORMANCE FOR EACH METRIC IS HIGHLIGHTED IN BOLD, AND THE SECOND-BEST IS UNDERLINED.

Method	Left hipbone				Left femur				Right femur			
	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)	Rot.(°)	Trans.(mm)	CD(mm)	PA(%)
Mean [45]	3.75	6.45	5.06	66.67	3.25	11.75	8.44	50.00	4.83	9.47	7.20	50.00
SSM [20]	<u>3.45</u>	5.90	4.23	66.67	<u>2.89</u>	11.89	7.67	50.00	<u>3.05</u>	9.99	7.64	50.00
AdaPoinTr [46]	3.73	**2.43**	<u>3.12</u>	**100.00**	4.75	<u>3.19</u>	<u>4.56</u>	**100.00**	5.71	**1.72**	**3.23**	**100.00**
FracFormer	**2.22**	<u>2.62</u>	**2.87**	**100.00**	**2.70**	**1.55**	**3.06**	**100.00**	**2.66**	<u>2.31</u>	<u>3.44</u>	**100.00**

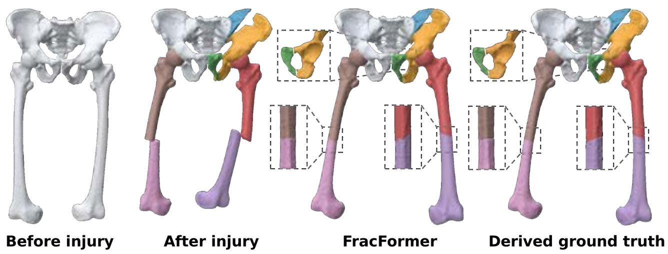

Fig. 14. Fractures and reduction planning results in the cadaver study.

difficult. Despite these challenges, our method ensured that small fragments remained within their anatomically plausible regions by leveraging global shape relationships and overall shape geometry. Additionally, errors in small fragments have minimal influence on the positioning of larger fragments, as demonstrated in Fig. 11, where the alignment of major fragments remains structurally stable and well-preserved.

F. Test on Special Cases

To further evaluate the robustness of our method, we tested its performance on three fractured cases that were not explicitly covered in the training or test dataset, including a case with incomplete CT scan, a pediatric fracture case, and a case with severe fracture deformity, which pose challenges for learning-based approaches.

As shown in Fig. 12, our method demonstrated strong generalization across these challenging scenarios. For fractures with incomplete CT scan, the method remained unaffected by missing regions due to its fragment-aware feature extraction, which enables it to handle incomplete anatomical structures. In the pediatric case, despite differences from adult anatomy, it achieved reasonable fragment realignment, likely benefiting from the model's ability to generalize across similar shapes and the separation of scale-related features in the DFGM. For severe deformity, the method provided a plausible reduction, though accuracy may decrease with extreme shape deviations.

G. Cadaver Study

We had the opportunity to conduct a cadaver study, where more reliable ground-truth poses free from annotation errors were derived using CT scan of intact bones prior to inducing fractures. A type C1.1 pelvic ring fracture was imparted to the cadaver's left hipbone, resulting in three fragments. And A3 transverse femoral shaft fractures were imparted to both femurs. Pre- and post-injury scans were both acquired with 0.98 mm in-plane pixel spacing and 0.90 mm slice thickness.

The results are summarized in Table VI and Fig. 14. Our method achieved clinically acceptable translational and rotational errors, with an average of 2.16 mm and 2.53°, respectively. Compared with real clinical fractures caused by high-energy trauma, these artificially induced fractures, created through manual bone chiseling, had simpler patterns. As a result, other methods also achieved competitive performance on some metrics. However, FracFormer consistently demonstrated the highest accuracy across all evaluated criteria.

V. DISCUSSION

A. Properties of the SSM

To evaluate the SSM's representation capacity, we assessed completeness, generality, and specificity following standard protocols [48]. Completeness measured the proportion of shape variance retained by the principal component modes, generality quantified reconstruction accuracy for unseen shapes via the Chamfer distance, and specificity evaluated the plausibility of randomly sampled shapes compared to the training data. The results were also summarized in Table I.

SSM effectively captured anatomical variability across different bones. For the hipbone and sacrum, which had larger sample sizes, we selected principal components that explained 95% of shape variability to reduce noise. In contrast, for the femur, where fewer samples were available, completeness was set to 99% to better preserve morphological differences. Under these completeness settings, generality remained below 2.32 mm, indicating accurate reconstructions, while specificity stayed within an acceptable range with a maximum of 4.15 mm, confirming that randomly sampled shapes were anatomically plausible. Together, these results show that the SSM captured diverse anatomical variations while ensuring accurate reconstructions and biological plausibility.

The relatively large number of principal components arose from removing scale differences before PCA, which necessitated more components to maintain completeness compared to models that retain size-related variations. This preprocessing focused the extracted modes purely on morphological differences, yielding a richer set of principal components. Fig. 10a illustrates the shape variations of the first three principal components without scale influence.

B. Properties of the Synthetic Fractures

We collected statistics on fracture fragment counts and sizes for the hipbone, sacrum, and femur across fracture

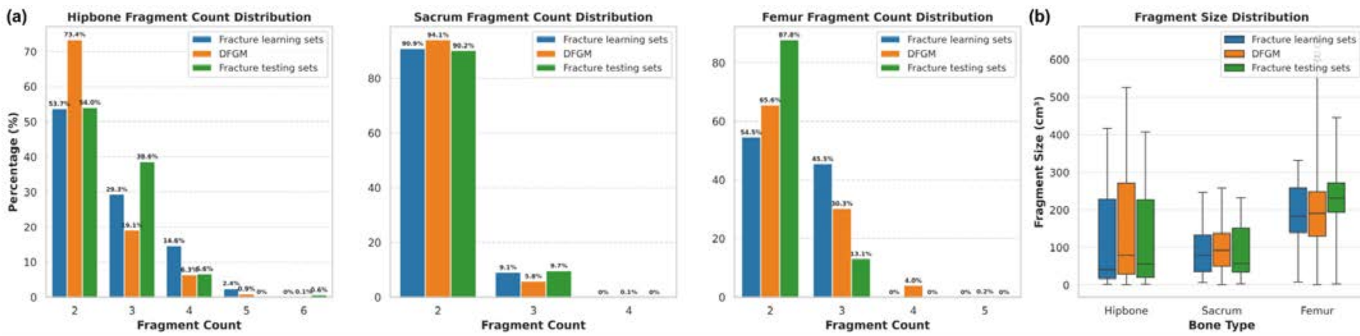

Fig. 15. Fragment count and size distributions across different datasets. (a) Histograms of fragment counts for hipbone, sacrum, and femur in fracture modes, DFGM-synthesized samples, and real test fractures. (b) Box plot of fragment size distributions across the three anatomies.

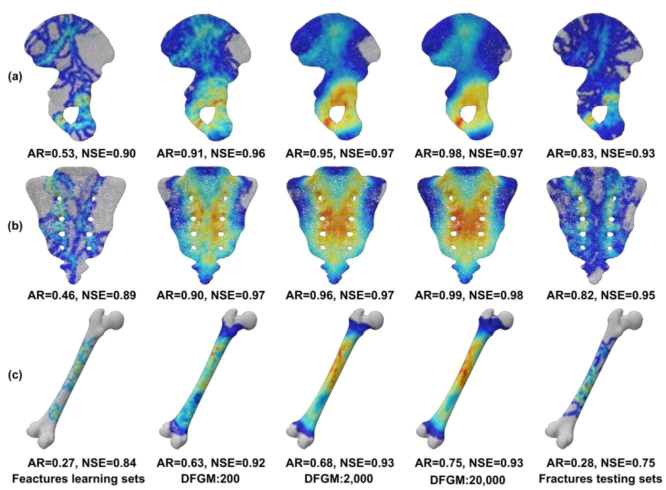

Fig. 16. Fracture distribution maps across different datasets. A blue-to-red gradient indicates increasing fracture occurrence, with non-fractured regions in gray. The corresponding AR and NSE values are annotated.

modes, DFGM samples, and real test fractures. Fig. 15a shows histograms of fragment counts, while Fig. 15b shows a box plot of fragment sizes. Compared to the limited variation in existing data, DFGM fractures exhibit a broader distribution of fragment counts and sizes, effectively covering the range observed in clinical test samples.

In addition, to assess the behavior of the synthesized fractures, we visualized the fracture distributions by computing fracture lines between adjacent fragments in each fractured sample and mapping them onto the SSM template. For each sample, we first derived its fracture mode representation on the SSM. The corresponding fracture lines were then extracted as a binary mask, where fractured regions were assigned 1 and non-fractured regions remained 0. To quantitatively assess the diversity of the fracture distribution, we employed two metrics: activation rate (AR) and normalized Shannon entropy (NSE). AR measures the proportion of template points fractured in at least one sample, reflecting the overall spread of fractures. NSE evaluates how evenly fractures are distributed across all points, where values closer to 1 indicate a broader, more uniform coverage of fracture lines.

Fig. 16 illustrates the fracture distribution maps across different datasets, along with the changes in fracture line patterns as the number of DFGM-generated samples increases from 200 to 20,000. The results indicate that both AR and NSE grow with more samples, signifying a wider spatial spread of fracture lines and enhanced fracture diversity. Starting from a limited set of fracture modes, DFGM generated more diverse fracture patterns, covering a larger anatomical region while maintaining clinically plausible distributions. Notably, high-frequency fracture regions aligned with clinical observations, such as frequent hipbone fractures occurring at the junction of the ilium, pubis, and ischium, and sacral fractures predominantly appearing around the structurally weaker sacral foramina.

C. Registration Approach

Our method falls into the category of template-based approaches, as it involves the prediction of the restored shape followed by the estimation of fragment poses. A notable advantage of our approach compared to other template-based methods is that the network output preserves patch-wise and fragment-wise correspondence with the input. This feature significantly simplifies and enhances the subsequent registration process. We have also tested fast point feature histogram (FPFH) for coarse registration, which has demonstrated effectiveness in prior studies [9], [14]. However, in our experiment, FPFH was outperformed by direct ICP, likely due to the more complex and unclear fragment geometry involved.

D. Limitations and Future Work

Due to the complexity of statistically modeling fracture patterns directly in a latent space, we instead map fractures from clinical cases onto simulated anatomies and introduce additional variations. Preliminary experiments have demonstrated the robustness and generalizability of this approach, showing that training with limited fracture modes is sufficient to infer unseen fracture patterns. However, for more irregular fracture patterns, additional experiments are required to assess the performance.

In our experiments, we focused on femoral shaft fractures, excluding proximal and distal fractures. This exclusion was due to the fact that CT scans for these regions are typically acquired locally, often missing the complete femoral shape. A promising approach to address this limitation would be to simulate data with intentionally excluded regions, enabling the method to better adapt to incomplete shape.

Another remaining challenge lies in handling small fragments, which possess limited features, making it difficult to capture their relationship to the holistic shape. Although fragments too small to be fixed are often ignored in clinical practice, those located in critical regions still warrant further investigation.

VI. Conclusion

We have introduced a novel integrated method for robust fracture reduction planning and realistic fracture simulation, addressing two fundamental challenges: the complexity of learning fracture-specific features and the scarcity of training data. The proposed network leverages fragment-aware patch encoding powered by transformer-based architectures, enabling accurate and generalizable fracture reduction for diverse fracture patterns observed in the hipbone, sacrum, and femur. By incorporating sparse-to-dense recursive refinement strategy, our method achieves precise alignment and robustness against diverse fracture configurations. Furthermore, the proposed DFGM effectively learns clinical fracture patterns from limited samples, generating diverse and realistic synthetic fractures. This sim-to-real approach enables training robust models that outperform those trained on real fracture data.

References

[1] M. J. Halawi, "Pelvic ring injuries: Surgical management and long-term outcomes," *J. Clin. Orthop. Trauma*, vol. 7, no. 1, pp. 1–6, 2016.

[2] P. R. Miller, P. S. Moore, E. Mansell, J. W. Meredith, and M. C. Chang, "External fixation or arteriogram in bleeding pelvic fracture: initial therapy guided by markers of arterial hemorrhage," *Journal of Trauma and Acute Care Surgery*, vol. 54, no. 3, pp. 437–443, 2003.

[3] M. A. Küper, A. Trulson, C. Minarski, F. Stuby, U. Stöckle, and C. Konrads, "Risks and strategies to avoid approach-related complications during operative treatment of pelvic ring or acetabular fractures," *Zeitschrift für Orthopädie und Unfallchirurgie*, vol. 159, no. 02, pp. 144–152, 2021.

[4] T. Tosounidis, N. Kanakaris, V. Nikolaou, B. Tan, and P. V. Giannoudis, "Assessment of lateral compression type 1 pelvic ring injuries by intraoperative manipulation: which fracture pattern is unstable?" *International orthopaedics*, vol. 36, pp. 2553–2558, 2012.

[5] C. Zhao, Q. Cao, X. Sun, X. Wu, G. Zhu, and Y. Wang, "Intelligent robot-assisted minimally invasive reduction system for reduction of unstable pelvic fractures," *Injury*, vol. 54, no. 2, pp. 604–614, 2023.

[6] M. Cimerman and A. Kristan, "Preoperative planning in pelvic and acetabular surgery: the value of advanced computerised planning modules," *Injury*, vol. 38, no. 4, pp. 442–449, 2007.

[7] A. S. Chowdhury, S. M. Bhandarkar, R. W. Robinson, and C. Y. Jack, "Virtual multi-fracture craniofacial reconstruction using computer vision and graph matching," *Computerized Medical Imaging and Graphics*, vol. 33, no. 5, pp. 333–342, 2009.

[8] E. M. Suero, T. Hüfner, T. Stübig, C. Krettek, and M. Citak, "Use of a virtual 3d software for planning of tibial plateau fracture reconstruction," *Injury*, vol. 41, no. 6, pp. 589–91, 2010.

[9] B. Zeng, H. Wang, X. Tao, H. Shi, L. Joskowicz, and X. Chen, "A bidirectional framework for fracture simulation and deformation-based restoration prediction in pelvic fracture surgical planning," *Medical Image Analysis*, vol. 97, p. 103267, 2024.

[10] M. Ulbrich, V. Van den Bosch, A. Bönsch, L. J. Gruber, M. Ooms, C. Melchior, I. Motmaen, C. Wilpert, A. Rashad, T. W. Kuhlen *et al.*, "Advantages of a training course for surgical planning in virtual reality for oral and maxillofacial surgery: crossover study," *JMIR Serious Games*, vol. 11, p. e40541, 2023.

[11] A. R. Willis, D. D. Anderson, T. P. Thomas, T. D. Brown, and J. L. Marsh, "3d reconstruction of highly fragmented bone fractures," in *SPIE Medical Imaging*, 2007.

[12] F. Paulano-Godino and J. J. Jiménez-Delgado, "Identification of fracture zones and its application in automatic bone fracture reduction," *Computer Methods and Programs in Biomedicine*, vol. 141, pp. 93–104, 2017.

[13] A. Luque-Luque, F. D. Pérez-Cano, and J. J. Jiménez-Delgado, "Complex fracture reduction by exact identification of the fracture zone," *Medical Image Analysis*, vol. 72, p. 102120, 2021.

[14] J. Liu, H. Li, B. Zeng, H. Wang, R. Kikinis, L. Joskowicz, and X. Chen, "An end-to-end geometry-based pipeline for automatic preoperative surgical planning of pelvic fracture reduction and fixation," *IEEE Transactions on Medical Imaging*, 2024.

[15] L.-h. Zhang, J.-x. Zhao, Z. Zhao, X.-y. Su, L.-c. Zhang, Y.-p. Zhao, and P.-f. Tang, "Computer-aided pelvic reduction frame for anatomical closed reduction of unstable pelvic fractures," *Journal of Orthopaedic Research*, vol. 34, no. 1, pp. 81–87, 2016.

[16] M. S. Ead, L. Westover, S. Polege, S. McClelland, J. L. Jaremko, and K. K. Duke, "Virtual reconstruction of unilateral pelvic fractures by using pelvic symmetry," *International Journal of Computer Assisted Radiology and Surgery*, vol. 15, pp. 1267–1277, 2020.

[17] C. Zhao, M. Guan, C. Shi, G. Zhu, X. Gao, X. Zhao, Y. Wang, and X. Wu, "Automatic reduction planning of pelvic fracture based on symmetry," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 10, no. 6, pp. 577–584, 2022.

[18] R. Gnat, E. Saulicz, M. Biały, and P. Kłaptocz, "Does pelvic asymmetry always mean pathology? analysis of mechanical factors leading to the asymmetry," *Journal of Human Kinetics*, vol. 21, no. 2009, pp. 23–32, 2009.

[19] R. Han, A. Uneri, R. C. Vijayan, P. Wu, P. Vagdargi, N. Sheth, S. Vogt, G. Kleinszig, G. M. Osgood, and J. H. Siewerdsen, "Fracture reduction planning and guidance in orthopaedic trauma surgery via multi-body image registration," *Med Image Anal*, vol. 68, p. 101917, 2021.

[20] S. Yibulayimu, Y. Liu, Y. Sang, G. Zhu, Y. Wang, J. Liu, C. Shi, C. Zhao, and X. Wu, "Pelvic fracture reduction planning based on morphable models and structural constraints," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interve.* Springer, 2023, pp. 322–332.

[21] S. Lu, Y. Yang, S. Li, L. Zhang, B. Shi, D. Zhang, B. Li, and Y. Hu, "Preoperative virtual reduction planning algorithm of fractured pelvis based on adaptive templates," *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 10, pp. 2943–2954, 2023.

[22] Y. Liu, S. Yibulayimu, Y. Sang, G. Zhu, C. Shi, C. Liang, Q. Cao, C. Zhao, X. Wu, and Y. Wang, "Preoperative fracture reduction planning for image-guided pelvic trauma surgery: A comprehensive pipeline with learning," *Medical Image Analysis*, p. 103506, 2025.

[23] Z. Deng, J. Jiang, Z. Chen, W. Zhang, Q. Yao, C. Song, Y. Sun, Z. Yang, S. Yan, Q. Huang *et al.*, "Tassembly: Data-driven fractured object assembly using a linear template model," *Computers & Graphics*, vol. 113, pp. 102–112, 2023.

[24] S. Yibulayimu, Y. Sang, Y. Liu, G. Zhu, Y. Wang, C. Zhao, Q. Cao, and X. Wu, "Automatic pelvic structure restoration: A sim-to-real approach via recursive pose estimation network," in *Proc. IEEE Int. Symp. Biomed. Imaging*. IEEE, 2024, pp. 1–5.

[25] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 652–660.

[26] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *ACM Transactions on Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.

[27] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, "Point transformer," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 16 259–16 268.

[28] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, "Pct: Point cloud transformer," *Computational Visual Media*, vol. 7, pp. 187–199, 2021.

[29] X. Yu, Y. Rao, Z. Wang, Z. Liu, J. Lu, and J. Zhou, "Pointr: Diverse point cloud completion with geometry-aware transformers," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 12 498–12 507.

[30] J. Adams and S. Elhabian, "Point2ssm: Learning morphological variations of anatomies from point cloud," *arXiv preprint arXiv:2305.14486*, 2023.

[31] D. Koschier, S. Lipponer, and J. Bender, "Adaptive tetrahedral meshes for brittle fracture simulation." in *Symposium on Computer Animation*, vol. 3, 2014.

[32] D. Hahn and C. Wojtan, "High-resolution brittle fracture simulation with boundary elements," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–12, 2015.

[33] J. Wolper, Y. Chen, M. Li, Y. Fang, Z. Qu, J. Lu, M. Cheng, and C. Jiang, "Anisompm: Animating anisotropic damage mechanics: Supplemental document," *ACM Trans. Graph*, vol. 39, no. 4, 2020.

[34] S. Raghavachary, "Fracture generation on polygonal meshes using voronoi polygons," in *ACM SIGGRAPH 2002 conference abstracts and applications*, 2002, pp. 187–187.

[35] S. Oh, S. Shin, and H. Jun, "Practical simulation of hierarchical brittle fracture," *Computer Animation and Virtual Worlds*, vol. 23, no. 3-4, pp. 291–300, 2012.

[36] Y.-C. Chen, H. Li, D. Turpin, A. Jacobson, and A. Garg, "Neural shape mating: Self-supervised object assembly with adversarial shape priors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 724–12 733.

[37] M. Müller, N. Chentanez, and T.-Y. Kim, "Real time dynamic fracture with volumetric approximate convex decompositions," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, pp. 1–10, 2013.

[38] J. Su, C. Schroeder, and R. Fedkiw, "Energy stability and fracture for frame rate rigid body simulations," in *Proc. 2009 ACM SIGGRAPH/Eurographics Symp. Comput. Anim.*, 2009, pp. 155–164.

[39] S. Sellán, J. Luong, L. Mattos Da Silva, A. Ramakrishnan, Y. Yang, and A. Jacobson, "Breaking good: Fracture modes for realtime destruction," *ACM Transactions on Graphics*, vol. 42, no. 1, pp. 1–12, 2023.

[40] Y. Liu, S. Yibulayimu, Y. Sang, G. Zhu, Y. Wang, C. Zhao, and X. Wu, "Pelvic fracture segmentation using a multi-scale distance-weighted neural network," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interve.* Springer, 2023, pp. 312–321.

[41] J. Adams and S. Elhabian, "Point2ssm++: Self-supervised learning of anatomical shape models from point clouds," *arXiv preprint arXiv:2405.09707*, 2024.

[42] T. Albrecht, M. Lüthi, T. Gerig, and T. Vetter, "Posterior shape models," *Med Image Anal*, vol. 17, no. 8, pp. 959–73, 2013.

[43] M. Attene, "A lightweight approach to repairing digitized polygon meshes," *The visual computer*, vol. 26, pp. 1393–1406, 2010.

[44] P. Liu, H. Han, Y. Du, H. Zhu, Y. Li, F. Gu, H. Xiao, J. Li, C. Zhao, L. Xiao *et al.*, "Deep learning to segment pelvic bones: large-scale ct datasets and baseline models," *International Journal of Computer Assisted Radiology and Surgery*, vol. 16, pp. 749–756, 2021.

[45] M. S. Ead, M. Palizi, J. L. Jaremko, L. Westover, and K. K. Duke, "Development and application of the average pelvic shape in virtual pelvic fracture reconstruction," *Int J Med Robot*, vol. 17, no. 2, p. e2199, 2021.

[46] X. Yu, Y. Rao, Z. Wang, J. Lu, and J. Zhou, "Adapointr: Diverse point cloud completion with adaptive geometry-aware transformers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 12, pp. 14 114–14 130, 2023.

[47] Y. Li, K. Mo, L. Shao, M. Sung, and L. Guibas, "Learning 3d part assembly from a single image," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*. Springer, 2020, pp. 664–682.

[48] B. C. Munsell, P. Dalal, and S. Wang, "Evaluating shape correspondence for statistical shape analysis: A benchmark study," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 11, pp. 2023–2039, 2008.