FRACTURE ASSEMBLY WITH SEGMENTATION AND ITERATIVE REGISTRATION

Jinhyeok Kim, Inha Lee, and Kyungdon Joo†
Artificial Intelligence Graduate School, UNIST
{jinhyeok,epsilon8854,kyungdon}@unist.ac.kr

ABSTRACT

Reassembling broken fractures back to their original shape remains a complex challenge. While prior research has demonstrated impressive results in domain-specific assembly, these methods largely depend on human-designed structural priors or struggle with assembling diverse shapes. To tackle this issue, we introduce a new fracture assembly framework based on segmentation and iterative registration, so-called FRASIER. By finding broken regions of fractures by segmentation, FRASIER dramatically increases overlap region ratios between fractures, which allows us to align fractures by registration. In addition, we employ point cloud XOR and beam search to make our framework robust. Experiments demonstrate that FRASIER outperforms state-of-the-art methods. Project page: https://frasier-assembly.github.io

Index Terms— fracture assembly, point cloud, segmentation, registration

1. INTRODUCTION

Fracture assembly (or shape assembly) is a reassembly task that pieces together the fragments of a broken object to restore its original shape [1, 2]; a jigsaw puzzle or the reassembling of broken skeletons or sculptures, for instance. Fracture assembly has gained attention in computer vision and robotics fields by virtue of its potential and usefulness across various applications, such as heritage archiving [3] and artifact preservation [4], and has remained a challenging task.

Fractures can be broadly classified into two types based on their causes: semantic and physical fractures. Concretely, one can decompose a given object into a set of segments according to their semantic meaning [5, 6]. This semantic fracture can be efficiently reassembled using both semantic and geometric cues. On the other hand, one can apply physical impacts on a given object, leading to more natural and diverse breaking patterns [4, 7, 8]. This physical fracture is more challenging than semantic fractures due to diverse and unexpected patterns and lack of semantic cues.

Let us imagine a scenario where we try to reassemble a set of physically broken fragments. We may first look at a set of broken fragments while picturing which fragments are geometrically adjacent at a high level. Then, we will choose two fragments that seem to fit together well among the broken fragments, attempting to assemble them into a single fragment by matching their adjacent sides. If it turns out that these two fragments were originally one piece, we will treat them as a single piece and proceed to select another fragment that is adjacent to this newly unified piece. By repeating this process, we can ultimately piece together the broken fragments into a single whole. This series of processes corresponds to an intuitive way that humans typically reassemble fractures.

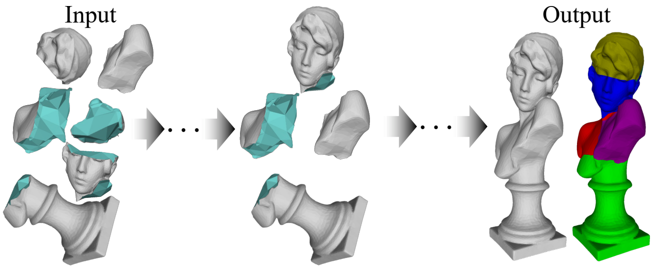

Fig. 1: FRASIER iteratively assembles fractures into the original shape, similar to how humans mend a fractured piece, thereby enhancing the robustness of our method.

Inspired by this intuitive way of humans, we aim to design a new framework for fracture assembly (see Fig. 1). Specifically, we employ a point cloud registration that aligns two sets of point clouds captured from different points of view [9, 10]. Nevertheless, it should be worth noticing that fracture assembly differs from the conventional point cloud registration in that broken fragments exhibit extreme conditions than scene registration data having enough overlap regions. To alleviate this gap, we present a fracture-specific registration approach that finds broken areas of fractures and then aligns them effectively. Moreover, we introduce the concept of point cloud XOR, which removes unnecessary points on broken regions after registration. Overall, our framework, FRASIER, recognizes broken regions of each fracture by point cloud segmentation, then iteratively aligns two feasible fractures by point cloud registration along with beam search to find diverse paths with trial and error. We evaluate the effectiveness of FRASIER on the Breaking Bad dataset [1]. Our framework shows robustness on fracture assembly tasks, largely outperforming previous state-of-the-art methods. Additionally, we thoroughly analyze each part of our framework to better understand its effect on the assembly outcome.

† Corresponding author

2. RELATED WORK

Automatic Part Assembly. Motivated by the desire to reconstruct 3D objects automatically, numerous approaches have been developed for the reassembly task. Some methods [11, 12, 13, 14] concentrate on semantic information from fractures gleaned from fracture patterns to accurately place fragments in their global positions. These approaches show effectiveness for reconstructing categorical objects that have commonly divisible parts, such as chairs, desks, and cars. In contrast, other methods [15, 16, 4] focus on geometrically divided fractures that exhibit randomness or lack semantic cues. To address the irregularity of fragments, these methods employ reconstruction priors [16] or local feature matching [15]. However, they still struggle with robustness when dealing with complexly broken objects in diverse categories. Our approach addresses this challenge by utilizing a powerful registration method, augmented by efficient point cloud removal and searching methods.

Point Cloud Registration. Point cloud registration is the task of aligning two point clouds that have a certain overlap, which can be applied to determine the relative pose of neighboring fractures. In the correspondence-based approach [9, 10, 17], correspondences between points in the two point clouds are first identified and serve as constraints that guide the alignment. With the advance of a data-driven approach, point features are obtained through an ML algorithm [17], or correspondence is directly obtained making the process end-to-end [10, 9]. For the alignment of fractures, we employ GeoTransformer [10] due to powerful tolerance on rotation and robust performance on low-overlap scenarios.

Point cloud segmentation. 3D point cloud segmentation can be applied to various scenarios, where point clusters exhibit shared intrinsic or semantic features. This versatility is enabled through permutation and translation invariance applied to model architecture. PointNet [18] introduced permutation invariance to manage irregular point formats via max pooling, while subsequent models [19, 20] leveraged features from local regions, akin to convolutional networks in 2D image processing. Especially, PointNext [21] shows superior performance and scalability which can serve as a semantic cue for broken surfaces, so that our method can utilize the semantic information even in complex fracture scenarios.

3. PROBLEM STATEMENT

Given a set of fractures $\mathcal{F} = \{\mathcal{P}_i\}$, where $\mathcal{P}_i$ denotes a point cloud uniformly sampled on the surface of the i-th fracture, fracture assembly aims to recover rotation $\mathbf{R}_i \in SO(3)$ and translation $\mathbf{t}_i \in \mathbb{R}^3$ of $\mathcal{P}_i$, thereby reconstructing original object shape. This problem can be considered a specific instance of a 3D registration problem. In particular, if we can distinguish two adjacent fractures that share a common surface among $\mathcal{F}$, we can reformulate this assembly problem into a set of registration problems in an iterative manner.

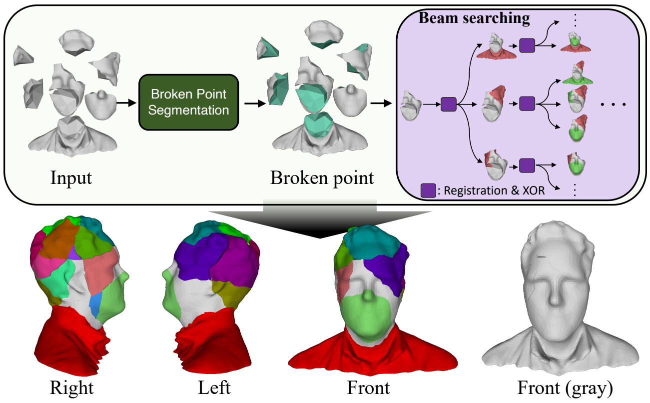

Fig. 2: **Overview of FRASIER**. Given fractures with its sampled point, we perform robust part assembly by utilizing point cloud segmentation, which allows us to alter the challenging part assembly problem into the classical point cloud registration. We visualize the point cloud as the shape for better understanding.

Let $\mathcal{P}_i$ and $\mathcal{P}_j$ be two fractures sharing a common surface, then assembly of $\mathcal{P}_i$ and $\mathcal{P}_j$ can be a registration problem focusing on finding relative pose for potential correspondences:

$$\underset{\mathbf{R},\mathbf{t}}{\arg\min} \sum_{(\mathbf{p},\mathbf{q})\in\mathbb{C}} \|(\mathbf{R}\mathbf{p}+\mathbf{t})-\mathbf{q}\|^2, \tag{1}$$

where $(\mathbf{p}, \mathbf{q})$ indicates a potential correspondence between $\mathcal{P}_i$ and $\mathcal{P}_j$ and $\mathbb{C}$ means the set of potential correspondences. However, unlike conventional registration problems with a certain overlap between point clouds, fractures have an extremely limited overlap, making this fracture assembly a non-trivial task. Additionally, since the fractures are randomly scattered, the connectivity between fractures remains uncertain, requiring the identification of adjacent pairs. Moreover, broken surfaces generally possess fewer unique features compared to scene registration data [22, 9], leading to inaccurate alignment of fractures.

4. METHOD

In this section, we present a new framework called FRASIER that gets inspiration from the intuitive way of human behavior for fracture assembly (see Fig. 2). Specifically, we first classify points on the broken surface region of each fracture, which allows us to focus on registration. We then iteratively perform a pair-wise point cloud registration for broken points; especially when we merge two fractures into a unified fracture, we exclude points on adjacent regions. In addition, we apply beam search within this fracture assembly framework to be robust against our pair-wise registration strategy.

4.1. Broken Point Segmentation and Registration

In the assembly process, broken regions often occupy a small portion of the overall surface area. The low overlap conditions pose difficulties for conventional registration methods, which are optimized for scenes with a substantial overlap ratio, ranging from 10-30% [9] to 30+% [22]. To address this issue, pinpointing these broken areas can assist in more accurate registration by eliminating unnecessary regions. For this purpose, we employ a point cloud segmentation to distinguish broken and unbroken regions of fractures. Specifically, we use Pointnext [21] for point cloud segmentation due to its ability to capture tiny local details as well as large-scale structures using multi-scale features. Based on Pointnext, we identify $\mathcal{P}$ of each fracture into broken segments $\mathcal{P}^b$ and non-broken one $\mathcal{P}^n$, where $\mathcal{P}$=$\mathcal{P}^b \cup \mathcal{P}^n$. We perform this broken point segmentation for every fracture before registration.

After individually segmenting the broken point of fractures, our next step is to find the relative pose among adjacent fractures. To this end, we employ a correspondence-based registration method based on Geotransformer [10], which shows robustness in low-overlap scenarios and is tolerant to rigid transformations. We first identify correspondences of fractures via Geotransformer, then apply the RANSAC [23] to filter out outliers. Subsequently, we determine rotation and translation using Singular Value Decomposition(SVD) and further refine them via Iterative Closest Point (ICP) [24].

4.2. Point cloud XOR

After registration, it is essential to combine the identified correct pairs of fragments. In our context, point clouds must be merged. The most basic approach is union operation on point cloud sets. However, union operation enlarges the size of the point cloud as the iteration progresses, increasing computational costs on later registrations. One way to prevent this is subsampling after the union, but this sacrifices registration accuracy as the point cloud becomes coarser.

To solve this problem, we introduce the concept of point cloud XOR. Point cloud XOR is a process of removing aligned points within a distance less than the threshold. This idea stems from the observation that broken regions overlap only once and therefore overlapped areas can be eliminated. This point cloud XOR scheme not only minimizes computation cost but also optimizes the overlap ratio between point clouds. Note that overlapped points belong to the broken points $\mathcal{P}^b$ where the unbroken points $\mathcal{P}^n$ have been pre-removed from the point cloud segmentation.

4.3. Beam Search

Even though segmentation and XOR modules greatly improve the overlap ratio between fractures, the registration results are not always accurate. Importantly, a single inaccurate result can adversely affect the overall reconstruction in an iterative process. To prevent this, we use beam search for diverse pathfinding, as outlined in [4]. Beam search is a search algorithm that explores a graph (align fractures) while maintaining a set of the most promising partial solutions observed so far. Algorithm details are explained below.

First, we choose n pairs of fractures, prioritizing those with more points. For these n pairs, we use a registration module to determine their relative poses. After identifying the poses, we merge the fractures using point cloud XOR. From the $n \cdot k$ results on pairs or n results at initial, select top k instances where the most points are removed via point cloud XOR. We iterate this process until a single object remains.

4.4. Implementation Detail

For the training segmentation model, we sparsely sample 50k points from each surface of fractures, and label which are classified as broken point cloud $\mathcal{P}^b$ or non-broken point cloud $\mathcal{P}^n$. During training, we used AdamW optimizer with cross-entropy loss between estimation and ground truth. For evaluation, we sample point cloud in proportion to the surface area of each fracture.

For the training registration model, we densely sampled points $\mathcal{P}^b$ from the broken surface of fractures as input and its corresponding $\mathbb{C}$ as labels. For inference, we use the estimated broken point $\hat{\mathcal{P}}^b$ from the segmentation module.

For beam search, we set $n = 6$ and $k = 3$. In other words, we choose 6 pairs from fractures with the largest area and evaluate 6 relative poses through the registration module. Given $3 \cdot 6$ results or 6 results initially, we choose the top 3 results and discard the rest.

5. EXPERIMENT

Dataset. We utilized Breaking Bad dataset [1], a physically simulated fracture dataset composed of three categories. Among three categories, we use `artifact` category data for training and `artifact,everyday` for evaluation.

Evaluation Metric. The metrics used for assessment include Root Mean Square Error (RMSE) for rotations and translations, as well as Part Accuracy (PA). PA is the percentage of parts perfectly positioned based on Chamfer Distance.

5.1. Evaluation

We evaluate the performance of `Frasier` in comparison to four baseline methods in the context of 3D fracture assembly tasks. All methods are trained on `artifact` category data. Qualitative results depicted in Fig. 3 provide a gallery of reconstruction from our framework and the baseline approaches. While other baselines struggle with simple fractures, our method can reconstruct from complex fractures. Quantitative results are presented in Table. 1. Our method

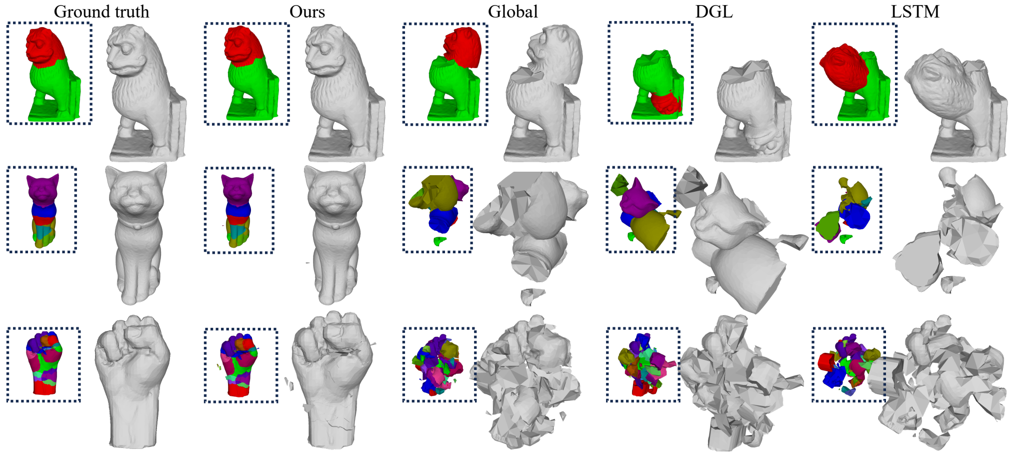

Fig. 3: We provide qualitative comparisons of `FRASIER` with Global [11, 12] , DGL [13], LSTM [14]. Compared to other methods, our approach exhibits superior quality in both simple (bottom) and complex (top) fractures.

Method	RMSE (R) ↓ degree	RMSE (T) ↓ $\times 10^{-2}$	PA ↑ %
Global [11, 12]	86.9	17.5	5.6
LSTM [14]	85.6	18.6	4.5
DGL [13]	86.3	18.0	9.6
Jigsaw [15]	52.4	22.2	45.6
Ours	**15.4**	**5.53**	**83.4**
Ours (**unseen**)	23.8	7.30	74.5

Table 1: **Quantitative comparison**. `FRASIER` outperforms previous methods in RMSE (R), RMSE (T), and PA. Ours (**unseen**) refers to results evaluated on unseen category data. Here, (R) and (T) denote rotation and translation, respectively.

significantly surpasses existing approaches, achieving a 30+ degree improvement in the RMSE(R) metric compared to Jigsaw [15]. In addition, we include evaluation results for unseen data in the final row of Table. 1. We evaluate on `everyday` to assess generalization performance. Our method, when tested on an unseen dataset, exhibits performance comparable to evaluations using the same category for training. This indicates our method is robust on distribution shift, which is essential to real-world application. Note that the visualization of Jigsaw [15] is excluded from Fig. 3 due to a lack of open-source code and Table. 1 directly quotes the quantitative result from Table.2 on [15].

5.2. Ablation study

We investigate the effect of individual modules on the quality of reconstruction. Fig. 4 presents ablation studies for broken point segmentation, point cloud XOR, and beam search. The top-left image is the outcome when all components are omitted, which is unsatisfactory as fractures often overlap with the volume or float in the air. However, as three components are added, the reconstruction becomes more realistic. Specifically, broken point segmentation and point cloud XOR greatly enhance the assembly process by eliminating non-essential points from fractures.

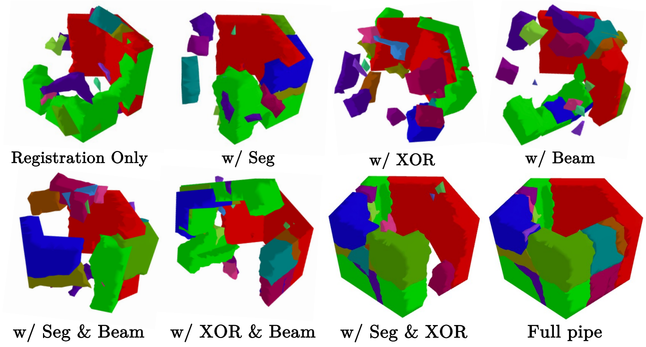

Fig. 4: **Ablation result from each module.** We compare results from simple registration to `FRASIER` by adding modules. Note that "Beam" refers to beam search, "Seg" refers to point cloud segmentation, and "XOR" refers to point cloud XOR.

6. CONCLUSION

In this paper, we present a novel framework that is inspired by how humans go through the process of assembling. By leveraging point cloud segmentation and registration methods for aligning, our framework can assemble complex fractures back to their original shape. In addition, experimental results show that point cloud XOR and beam search substantially improve registration quality. Finally, `FRASIER` significantly outperforms state-of-the-art methods and demonstrates robustness when evaluated on unseen data during training.

7. ACKNOWLEDGEMENT

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI and No.2020-0-01336, Artificial Intelligence Graduate School Program (UNIST)), and Artificial intelligence industrial convergence cluster development project funded by the Ministry of Science and ICT (MSIT, Korea) & Gwangju Metropolitan.

8. REFERENCES

[1] Silvia Sellán, Yun-Chun Chen, Ziyi Wu, Animesh Garg, and Alec Jacobson, "Breaking bad: A dataset for geometric fracture and reassembly," *NeurIPS*, 2022.

[2] Thomas Funkhouser, Hijung Shin, Corey Toler-Franklin, Antonio García Castañeda, Benedict Brown, David Dobkin, Szymon Rusinkiewicz, and Tim Weyrich, "Learning how to match fresco fragments," *JOCCH*, 2011.

[3] Ruggero Pintus, Kazim Pal, Ying Yang, Tim Weyrich, Enrico Gobbetti, and Holly Rushmeier, "A survey of geometric analysis in cultural heritage," *Computer Graphics Forum*, 2016.

[4] Je Hyeong Hong, Seong Jong Yoo, Muhammad Arshad Zeeshan, Young Min Kim, and Jinwook Kim, "Structure-from-sherds: Incremental 3d reassembly of axially symmetric pots from unordered and mixed fragment collections," *CVPR*, 2021.

[5] Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su, "PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding," *CVPR*, 2019.

[6] Benjamin Jones, Dalton Hildreth, Duowen Chen, Ilya Baran, Vladimir G Kim, and Adriana Schulz, "Automate: A dataset and learning approach for automatic mating of cad assemblies," *ACM TOG*, vol. 40, no. 6, pp. 1–18, 2021.

[7] Nikolas Lamb, Cameron Palmer, Benjamin Molloy, Sean Banerjee, and Natasha Kholgade Banerjee, "Fantastic breaks: A dataset of paired 3d scans of real-world broken objects and their complete counterparts," *CVPR*, 2023.

[8] Silvia Sellán, Jack Luong, Leticia Mattos Da Silva, Aravind Ramakrishnan, Yuchuan Yang, and Alec Jacobson, "Breaking good: Fracture modes for realtime destruction," *ACM TOG*, vol. 42, no. 1, 2023.

[9] Shengyu Huang, Zan Gojcic, Mikhail Usvyatsov, Andreas Wieser, and Konrad Schindler, "Predator: Registration of 3d point clouds with low overlap," *CVPR*, 2021.

[10] Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, and Kai Xu, "Geometric transformer for fast and robust point cloud registration," *CVPR*, 2022.

[11] Nadav Schor, Oren Katzir, Hao Zhang, and Daniel Cohen-Or, "Componet: Learning to generate the unseen by part synthesis and composition," *CVPR*, 2019.

[12] Jun Li, Chengjie Niu, and Kai Xu, "Learning part generation and assembly for structure-aware shape synthesis," *AAAI*, 2020.

[13] Guanqi Zhan, Qingnan Fan, Kaichun Mo, Lin Shao, Baoquan Chen, Leonidas J Guibas, Hao Dong, et al., "Generative 3d part assembly via dynamic graph learning," *NeurIPS*, 2020.

[14] Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen, "Pq-net: A generative part seq2seq network for 3d shapes," *CVPR*, 2020.

[15] Jiaxin Lu, Yifan Sun, and Qixing Huang, "Jigsaw: Learning to assemble multiple fractured objects," *arXiv preprint arXiv:2305.17975*, 2023.

[16] Yun-Chun Chen, Haoda Li, Dylan Turpin, Alec Jacobson, and Animesh Garg, "Neural shape mating: Self-supervised object assembly with adversarial shape priors," *CVPR*, 2022.

[17] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich, "Superpoint: Self-supervised interest point detection and description," *CVPR workshops*, 2018.

[18] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," *CVPR*, 2017.

[19] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *NeurIPS*, 2017.

[20] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen, "Pointcnn: Convolution on x-transformed points," *NeurIPS*, 2018.

[21] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem, "Pointnext: Revisiting pointnet++ with improved training and scaling strategies," *NeurIPS*, 2022.

[22] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser, "3dmatch: Learning local geometric descriptors from rgb-d reconstructions," *CVPR*, 2017.

[23] Martin A Fischler and Robert C Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[24] Szymon Rusinkiewicz and Marc Levoy, "Efficient variants of the icp algorithm," *Proceedings third international conference on 3-D digital imaging and modeling*, pp. 145–152, 2001.