

Scalable Geometric Fracture Assembly via Co-creation Space among Assemblers

Ruiyuan Zhang¹*, Jiayang Liu¹*, Zexi Li¹, Hao Dong³, Jie Fu²*[†], Chao Wu¹[†]

¹ Zhejiang University

² Hong Kong University of Science and Technology

³ Peking University

{zhangruiyuan,zjljx,zexi.li,chao.wu}@zju.edu.cn, jjiefu@ust.hk, hao.dong@pku.edu.cn

Abstract

Geometric fracture assembly presents a challenging practical task in archaeology and 3D computer vision. Previous methods have focused solely on assembling fragments based on semantic information, which has limited the quantity of objects that can be effectively assembled. Therefore, there is a need to develop a scalable framework for geometric fracture assembly without relying on semantic information. To improve the effectiveness of assembling geometric fractures without semantic information, we propose a co-creation space comprising several assemblers capable of gradually and unambiguously assembling fractures. Additionally, we introduce a novel loss function, i.e., the geometric-based collision loss, to address collision issues during the fracture assembly process and enhance the results. Our framework exhibits better performance on both PartNet and Breaking Bad datasets compared to existing state-of-the-art frameworks. Extensive experiments and quantitative comparisons demonstrate the effectiveness of our proposed framework, which features linear computational complexity, enhanced abstraction, and improved generalization. Our code is publicly available at <https://github.com/Ruiyuan-Zhang/CCS>.

Introduction

Fracture assembly aims to reconstruct a broken object by composing its fractures. Manually assembling these fragments is time-consuming and requires precision. The task is complicated due to the vast number of potential combinations and the lack of clear instructions. Therefore, geometric fracture assembly is a practical but challenging task.

Previous studies have focused on tasks like archaeological fragment matching and 3D furniture assembly, using methods such as generative 3D part assembly with graph neural networks to understand part relationships (Funkhouser et al. 2011; Toler-Franklin et al. 2010; Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022; Lee, Hu, and Lim 2021; Li et al. 2020). However, these approaches often face limitations in handling shapes with numerous fragments, also called scalability issue or multi-parts issue. Furthermore, they rely heavily on semantic information like

*These authors contributed equally.

[†]Corresponding authors.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Comparison of Transformer-based assembly (left) and our method (right). Left: Assemblers may become confused when presented with multiple fractures to assemble simultaneously, leading to an inability to assemble any parts. Right: In our method, assemblers focus on essential fractures of parts, enabling efficient and effective task completion.

part segmentation and ground-truth annotations (Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022; Lee, Hu, and Lim 2021; Li et al. 2020). To address part relationship constraints, (Zhang et al. 2022) introduced a transformer-based framework, IET, which assigns positions to parts and uses self-attention for positional interactions. However, this method’s computational complexity increases significantly with the number of parts due to the quadratic scaling of attention mechanisms (Goyal et al. 2021).

In this paper, we introduce the Co-creation Space, a novel approach that enables assemblers to predict the 6-DoF pose of geometric parts or fractures. This method, illustrated in Figure 2, involves assemblers competing for write access to update a shared workspace, thereby predicting part poses (Baars 1993; Dehaene, Lau, and Kouider 2021). The task of geometric fracture assembly is divided into specialized tasks, such as identifying the core fracture and locating relevant ones. Assemblers collaborate in a shared workspace, which replaces the pairwise interactions found in traditional dot-product attention methods. This approach ensures global coherence and reduces computational complexity to a linear scale relative to the number of assemblers (Baars 1993). As depicted in Figure 1, the Co-creation Space also acts as an information bottleneck, denoted as R_t in Eq 2, limiting the capacity of information channels between specialists. This ensures that only essential information is written into the workspace. Fig.1(a): when the assembler considers multiple parts/fractures simultaneously, it discovers that multiple

components can match each other (e.g., a chair seat can connect with both the chair legs and the chair back; a fragment can link with other fragments at multiple angles). As the volume of information increases, it would cause confusion issue as other baselines (Fig.1(b)). Just as drivers pay more attention to important traffic signals rather than the onboard music while driving, the assembler requires a focus.

Moreover, the presence of identical or similar information in fractures during assembly can cause ambiguity and hinder escaping local optima. We address this by introducing a geometric-based collision loss, depicted in Figure 3. This loss function actively separates identical or similar fractures that share the same location, guiding them towards more appropriate positions. We carried out comprehensive experiments on two major geometric fracture assembly datasets: PartNet (Mo et al. 2019) and Breaking Bad (Sellán et al. 2022). Through numerous comparative experiments and ablation analyses, we compared our method with state-of-the-art works (Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022) and verified the effectiveness of our proposed framework. Our method emphasizes scalability in assembly processes and can be integrated as a plug-and-play module to enhance geometric information extraction research.

Related Works

Geometric Shape Assembly

Geometric shape assembly involves combining multiple shapes to create a target object (Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022; Grason 2016; Chen et al. 2022; Funkhouser et al. 2011; Jones et al. 2021; Lee, Hu, and Lim 2021; Li et al. 2020; Liu et al. 2023a), and it is an important problem in science and engineering (Funkhouser et al. 2011; Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022). Previous research has focused on specific cases that simplify the problem, such as using identical fragments (Funkhouser et al. 2011; Lee et al. 2022) or textured fragments (Lee et al.). However, in practical applications, the shapes and number of fractures can be arbitrary, requiring more general methods.

Building on PartNet (Mo et al. 2019), studies like (Zhan et al. 2020; Narayan, Nagar, and Raman 2022) have proposed graph-based learning methods for predicting 6-DoF poses of each part and assembling a shape. Similarly, (Zhang et al. 2022) use a transformer-encoder to understand part relationships. However, these methods often rely heavily on semantic information of object parts, such as instance encoding (Zhang et al. 2022), instance label (Zhan et al. 2020), and order information (Narayan, Nagar, and Raman 2022), and become less effective with more parts. Recently, (Sellán et al. 2021) introduced the Breaking Bad dataset, highlighting the challenge of assembling non-semantic fractures into complete shapes. This shows that fractured shape assembly is an ongoing issue. In response, we propose a new framework to tackle the scalability challenge in geometric shape assembly, demonstrating good performance on non-semantic datasets.

Shared Global Workspace

In cognitive neuroscience, the Global Workspace Theory (SGW) (Baars 1993; Dehaene, Lau, and Kouider 2021) has been proposed to suggest an architecture allowing specialist modules to interact through a shared representation called workspace, a bandwidth-limited communication channel. The advantage of this approach is that it enables global coordination and coherence among different components, beyond just local or pairwise interactions. Workspace can be modified by any specialist, and that is broadcast to all specialists. (Goyal et al. 2021) explore the use of such a communication channel in the context of deep learning, leading to greater abstraction and better generalization, which will be effective in large geometric fracture assembly tasks without semantic information. Based on this theory, (Liu et al. 2022) propose a Centralized Training Decentralized Execution learning approach called Stateful Active Facilitator that enables agents to work efficiently in high-coordination and high-heterogeneity environments. In this paper, co-creation is proposed so that the assemblers competing for write access can update the workspace. Looking at it from another perspective, the process in SGW is similar to an information bottleneck (IB) that distills and extracts crucial information for our conscious awareness. In this paper, IB corresponds to R_t , used to address scalability issues.

Method

Let $\mathcal{P} = \{p_i\}_{i=1}^N$ represent a set of geometric fracture point clouds, where $p_i \in \mathbb{R}^{n_{pc} \times 3}$ and N denotes the number of parts which may vary for different 3D shapes. Our goal is to predict a set of 6-DoF fracture poses as $Z = \{(R_i, T_i)\}_{i=1}^N$ in $SE(3)$ space, where $R_i \in \mathbb{R}^4$ and $T_i \in \mathbb{R}^3$ denotes the rigid rotation and translation for each fracture, respectively. Then, we apply the predicted pose to transform the point cloud of each part and get the i -th fracture’s predicted point cloud $\mathcal{P}'_i = Z_i(p_i)$, in which Z_i is the joint transformation of (R_i, T_i) . And the complete shape can be assembled into $S = \bigcup_{i=1}^N \mathcal{P}'_i$ as our predicted assembly result. In this work, we present a scalable geometric fractures assembly framework via a co-creation space among assemblers to assemble 3D shapes, which is illustrated in Figure 2.

Co-creation Space among Assemblers

In this section, we introduce co-creation space among assemblers, which serves as the core module of this framework. In this module at each computation stage indexed by t , n_a assemblers compete for write access to co-creation space, $n_a = N$. The contents of the co-creation space, in turn, are broadcast to all assemblers simultaneously.

Step 1: Generating the message of assemblers. Each assembler i receives a message containing geometric fractures’ information $m_{i,t}$ at each time step t . The initialization information is created through the routing function, which carries information about the current fracture and its relationship with other fractures. The first step is external to the co-creation space, for details, please refer to . we denote the set of messages generated by the assemblers at time step t




Figure 2: The overall architecture of our framework. (a) geometric fractures feature extraction using shared parameters PointNet (Qi et al. 2017), (b) geometric fractures relational reasoning and information routing, (c) complete the assembly task among assemblers using the co-creation space, (d) MLP predictor for pose estimation.

by M_t :

$$M_t = \{m_{i,t} | 1 \leq i \leq N\}. \quad (1)$$

Step 2: Writing into the co-creation space. The message M_t generated in step one is distilled into a latent state which we term as a **Co-creation Space**. The assemblers compete to write into the co-creation space, whose contents need to be updated in the context of new information received from different assemblers. This step ensures that only the critically important signals make it to the co-creation space, therefore preventing the co-creation space from being cluttered. We represent the co-creation space state at time step t by R_t . R_t consists of L slots $\{l_0, l_1, \dots, l_{L-1}\}$, each of dimension d_l so that $R_t \in \mathbb{R}^{L \times d_l}$. The messages in M_t compete to write into each co-creation space's memory slot via a key-query-value attention mechanism. In this case, the query is a linear projection of the state of the current co-creation space memory content R_t , i.e., $\tilde{Q} = R_t \tilde{W}^q$, whereas the keys and values are linear projections of the messages M_t . Co-creation space state is updated as:

$$R_t = \text{softmax}\left(\frac{\tilde{Q}(M_t \tilde{W}^e)^T}{\sqrt{d_e}}\right) M_t \tilde{W}^v. \quad (2)$$

In this work, we use a top- k softmax (Ke et al. 2018) to select a fixed number of assemblers allowed to write in the co-creation space. Similar to transformer (Vaswani et al. 2017), we apply multiple heads to improve its expressive ability.

Step 3: Reading from the co-creation space. Each assembler then updates its state using the information broadcast from the co-creation space. We again utilize an attention mechanism to perform this operation. All the assemblers create queries $Q_t = \{q_{i,t} | 1 \leq i \leq N\} \in \mathbb{R}^{N \times d_e}$ where

$q_{i,t} = W_{read}^q a_{i,t}$ and $a_{i,t}$ are encoded partial observations of one assembler. Generated queries are matched with the keys $K = R_t W^e \in \mathbb{R}^{L \times d_e}$ from the updated memory slots of co-creation space. As a result, the attention mechanism can be written as:

$$M'_t = \text{softmax}\left(\frac{Q_t K^T}{\sqrt{d_e}}\right) R_t W^v, \quad (3)$$

where $M'_t = \{m'_{i,t} | 1 \leq i \leq N\}$. After receiving the broadcast information from the co-creation layer, each assembler updates its state by a feedforward layer. This yields the new value M_{t+1} for the k -th assembler, from which we start the next stage ($t+1$).

Co-creation is a shared workspace with limited capacity, which encourages specialization and compositionality among assemblers. IET (Zhang et al. 2022) relies on pairwise interactions captured via an attention mechanism. Unfortunately, such attention mechanisms scale quadratically with the number of assemblers. Here, the computational complexity of the proposed method is linear in the number of assemblers.

Geometric Fractures Information Routing

Inspired by modular deep learning (Pfeiffer et al. 2023), before the parts assembly task, we employed an independent module to perform geometric information extraction and fractures relation reasoning, and route the aggregation information to the next module. A transformer-based architecture is recommended to learn the relationships between fractures. Transformer-encoder applies multiple self-attention layers (Zhu et al. 2023a,b) that aggregate information (e.g., geometry and posture) from the entire input sequence (here is a set of geometric fractures). The positional encoding is

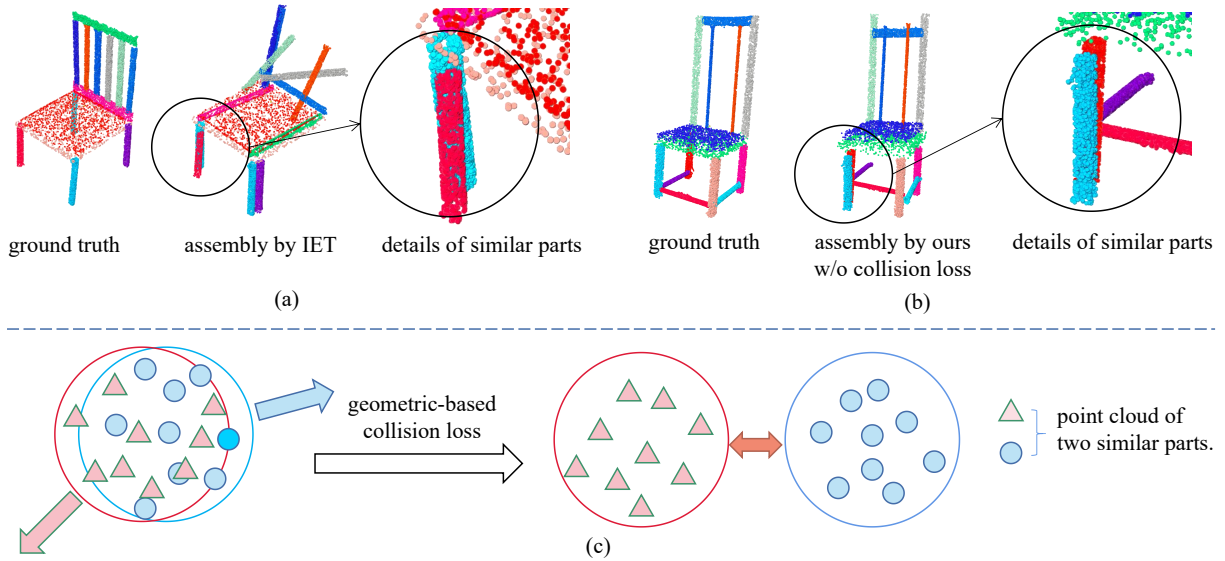


Figure 3: The illustration of geometric-based collision loss (The red circles and the blue triangles represent 2 similar fractures).

omitted as the input already contains the information about 3-dimensional XYZ coordinates. We didn't follow the standard formulation of the transformer. Instead, in order to keep the model architecture concise, this framework used the same Transformer structure as the co-creation space. Here the co-creation space can support information routing and relation reasoning at structure, but these two modules are different in terms of their functions in our framework.

The module discussed mainly performs information extraction and routing, while the co-creation space primarily completes the logical task of fracture assembly. Learning routing is a challenging research of modular deep learning, including *training instability* (Pfeiffer et al. 2023), *module collapse* (Kirsch, Kunze, and Barber 2018), and *overfitting* (Pfeiffer et al. 2023). This is not the focus of this paper to be discussed, we will discuss the learned routing of Assembly in future work. We have provided several options to replace the geometric fracture information routing network structure, including ResNet (He et al. 2016), Transformer (Vaswani et al. 2017), and GNN (Scarselli et al. 2008).

Geometric-based Collision Loss

During assembling, fractures with similar or identical semantic information will be placed in the same location. As shown in Figure 3, other chair legs or similar parts are assembled in the same position. To solve this problem, previous methods manually input the semantic information of the parts, such as Instance Encoding (Zhang et al. 2022), Instance label (Zhan et al. 2020), and order information (Narayan, Nagar, and Raman 2022). However, these methods are not suitable for assembly tasks that do not involve semantic information. Additionally, adding too much bias into the model will limit the generalization ability of these methods to new shape distributions.

We hypothesize that the problem arises because of *frac-*

tures may assemble at a local optimum point, which is extremely difficult to escape. When one of the fractures is gently moved, the model will still pull the moved part back to its local optimum position during optimization. To verify our hypothesis, we propose a loss function called *collision loss*, which warns and pushes away one of the fractures when placed in the same location. The displaced fracture will affect other losses, such as the shape chamfer distance loss, and gradually optimize it to a reasonable position. Experimental results show that our method effectively solves the ambiguity between fractures c_i , where $1 \leq i \leq N$. we define the Collision loss as:

$$\mathcal{L}_c = \frac{2 \times \sum_i^N \sum_{j < i}^N (1 - \log(C \|c_i - c_j\|_2))}{N \times (N - 1)}, \quad (4)$$

where C is the hyperparameter of collision loss, that can be adjusted through grid search (Liu et al. 2023b). Which presents a correlation with the distance to the cloud point distribution. $\|c_i - c_j\|_2$ represents the distance $d \geq 0$ between two parts or fractures. We only consider the loss $l_{i,j} = 1 - \log(C \times d)$. The corresponding derivative is $l'_{i,j} = -\frac{1}{d}$. As d increases, $l'_{i,j}$ becomes larger, meaning its absolute value decreases because the derivative is negative. The increase of the derivative leads to the slope of the $L_{i,j}$ becoming flatter. It indicates that the rate of decrease of the function slows down as d increases. This characteristic explains that \mathcal{L}_c can timely separate two overlapping fractures without affecting the total loss of two non-overlapping parts. This can effectively solve the ambiguous issue. We add an ablation study for C with artifact dataset in Table 2.

Training Details

In the task of geometric fracture assembly, there are multiple possible solutions due to the interchangeable positions of the

fractures and the ability to place decorative parts in different locations. In order to establish an uncertainty model and explore structural diversity, the Min-of-N (MoN) loss (Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022) and random noise $z_j \sim \mathcal{N}(0, 1)$ were adopted. Here the overall framework was defined as f , while the ground truth pose was defined as f^* . The MoN loss is used to calculate the error, which is defined as follows:

$$\mathcal{L}_{MoN} = \min_{1 \leq j \leq n} \mathcal{L}(f(\mathcal{P}, z_j), f^*(\mathcal{P})). \quad (5)$$

Given a set of fracture point clouds \mathcal{P} , f make n predictions by perturbing the input with n random vector z_j . Intuitively, it ensures at least one prediction as close as the ground-truth space. Following (Zhang et al. 2022; Zhan et al. 2020; Narayan, Nagar, and Raman 2022), we set $n = 5$ in the experiment. The loss function, \mathcal{L} is split into four categories similar to (Zhan et al. 2020), collision loss was proposed by this work, for global and part-wise structural integrity.

The translation is supervised by Euclidean loss \mathcal{L}_t , which measures the distance between the predicted translation T_i and ground-truth translation T_i^* for each fracture,

$$\mathcal{L}_t = \sum_{i=1}^N \|T_i - T_i^*\|_2^2. \quad (6)$$

The rotation is supervised via Chamfer distance on the rotated fracture point cloud:

$$\mathcal{L}_r = \sum_{i=1}^N \left(\sum_{x \in R_i(p_i)} \min_{y \in R_i^*(p_i)} \|x - y\|_2^2 + \sum_{x \in R_i^*(p_i)} \min_{y \in R_i(p_i)} \|x - y\|_2^2 \right), \quad (7)$$

in which the $R_i(p_i)$ and $R_i^*(p_i)$ represent the rotated fracture points p_i using the estimated rotation R_i and the ground-truth R_i^* , respectively.

To ensure comprehensive assembly quality, we have employed Chamfer distance (CD) to monitor the entire shape assembly S .

$$\mathcal{L}_s = \sum_{x \in S} \min_{y \in S^*} \|x - y\|_2^2 + \sum_{y \in S^*} \min_{x \in S} \|x - y\|_2^2, \quad (8)$$

where S is the assembled shape and S^* denotes the ground-truth. The total loss is defined as follows:

$$\mathcal{L} = w_c \mathcal{L}_c + w_t \mathcal{L}_t + w_r \mathcal{L}_r + w_s \mathcal{L}_s, \quad (9)$$

where w_c , w_t , w_r and w_s denote the weight of different losses, which are empirically determined.

Experiments and Analysis

Dataset and Baselines

We evaluated our method and baselines on PartNet (Mo et al. 2019) datasets and Breaking Bad (Sellán et al. 2022, 2021). PartNet is a large-scale shape dataset with fine-grained and

hierarchical part segmentation. We utilize the Chair dataset with default train/validation/test being split in the dataset, which includes 6,323 chairs. Breaking Bad dataset (Sellán et al. 2021) contains a diverse set of shapes spanning everyday objects, artifacts without any manually annotated semantic information, e.g., instance label. It combines one million geometrically natural fracture patterns, which meets our needs. We used the categories of Everyday and Artifact. The number of parts or fractures used in all datasets ranges from 2 to 20. We compare the proposed method with Global (Li et al. 2020; Schor et al. 2019), LSTM (Wu et al. 2020a), CM (Sung et al. 2017), DGL (Zhan et al. 2020), RGL (Narayan, Nagar, and Raman 2022), IET (Zhang et al. 2022). To ensure a fair comparison, the implementations of other baselines are based on benchmark (Sellán et al. 2022).

Evaluation Metrics

The performance of our proposed method and baselines are measured by generating a variety of shapes and finding the closest shape to the ground truth using minimum matching distance (Achlioptas et al. 2018). To ensure a thorough evaluation, we use the metrics of part accuracy (PA) (Li et al. 2020), connectivity accuracy (CA) (Zhan et al. 2020), and shape Chamfer distance (SCD) (Zhan et al. 2020), as employed by (Zhan et al. 2020). PA and CA assess the precision of each individual part and the quality of the connections between them, respectively, while SCD evaluates the overall quality of the assembled shape. In addition, we computed the root mean squared error RMSE(R) (Sellán et al. 2022) and root mean squared error RMSE(T) (Sellán et al. 2022) to evaluate both rotation and translation prediction when conducting experiments using the Breaking Bad dataset.

Experimental Results and Analysis

We conducted a performance evaluation of our proposed method and various baselines, as illustrated in Table 4. Our proposed method demonstrated superior performance in most columns, particularly in the part and connectivity accuracy metrics. The red numerical value represents the number of instances where the proposed method outperforms the second-best method, and the green numerical value signifies the extent by which the proposed method lags behind the best-performing method.

The visual outcomes depicted in Figure 4 and Figure B.2, Figure B.3 provide empirical evidence that our proposed approach surpasses the baseline methods in generating meticulously organized geometries. In contrast, the baseline methods often fail to achieve satisfactory assembly results. Subjectively speaking, our framework shows results that are almost indistinguishable from the ground truth. Figure 4 [d-h] shows that our framework can perform well even with many parts. This is because our framework includes a co-creation space, which has an information bottleneck property that promotes the emergence of better assembler professional skills, resulting in better generalization and logical reasoning abilities for the network (Tishby, Pereira, and Bialek 2000; Tishby and Zaslavsky 2015; Wu et al. 2020b; Ahuja et al. 2021; Goyal et al. 2021). From Figure 4 [c-h],

Metrics	Category	LSTM	Global	CM	DGL	RGL	IET	Ours
SCD ($\times 10^{-3}$) \downarrow	Chair	13.10	14.60	24.10	09.10	08.70	09.40	07.00
	Everyday	20.42	-	-	15.12	-	14.29	14.26
	Artifact	25.42	-	-	17.47	-	16.46	15.49
PA \uparrow	Chair	21.77	15.70	08.78	39.00	49.06	37.50	53.59
	Everyday	18.30	-	-	26.40	-	26.94	28.57
	Artifact	06.22	-	-	16.09	-	17.64	19.78
CA \uparrow	Chair	06.89	09.90	09.19	23.87	32.26	24.47	38.97
RMSE (R) \downarrow	Everyday	83.50	-	-	80.99	-	80.49	79.18
	Artifact	86.20	-	-	82.64	-	79.94	77.59
RMSE (T) ($\times 10^{-2}$) \downarrow	Everyday	16.63	-	-	15.35	-	15.07	15.10
	Artifact	17.50	-	-	16.05	-	15.72	15.50

Table 1: Quantitative evaluation on PartNet (Chair) and Breaking Bad datasets (Everyday and Artifact).

w_c	C	SCD $\times 10^{-3}$ \downarrow	PA \uparrow	RSME $_T$ \downarrow	RSME $_R$ \downarrow
0.10	5	1.717	19.53	0.1618	80.686
0.10	10	1.610	19.81	0.1585	81.049
0.10	15	1.651	19.10	0.1589	80.854
0.10	20	1.615	19.83	0.1566	80.054
0.10	25	1.603	19.83	0.1560	79.977
0.10	30	1.581	20.53	0.1554	80.811
0.10	35	1.623	19.39	0.1562	80.509
0	N/A	1.631	19.81	0.1579	80.507
0.05	30	1.685	19.10	0.1565	81.028
0.15	30	1.587	20.54	0.1553	80.261
0.20	30	1.645	19.74	0.1573	79.837
0.25	30	1.710	19.38	0.1561	80.442

Table 2: The grid search of w_c and C in artifact.

Workspace	PA \uparrow	SCD($\times 10^{-3}$) \downarrow	CA \uparrow
w/o any workspace	45.89	08.10	31.73
$k = 1$ w/o Col. loss	49.87	07.60	35.65
$k = 5$ w/o Col. loss	49.19	07.30	33.75
$k = 10$ w/o Col. loss	<u>52.53</u>	<u>07.10</u>	<u>38.40</u>
$k = 15$ w/o Col. loss	52.52	07.32	36.90
$k = 20$ w/o Col. loss	50.23	07.30	35.56
$k = 10$ with Col. loss	53.59	07.00	38.97

Table 3: The grid search of k in our settings and ablation study. ‘‘Col. loss’’ denotes collision loss.

we can see that in the baseline, similar part problems cannot be handled well when collisions occur. However, this problem is resolved in our framework. The success of our framework relies on the proposed collision loss, which can identify local optimal points during model training and overcome this state, and lead better results. Figure 4 [a-b] shows that our framework still performs well on Breaking Bad, demonstrating that our framework has better relational reasoning abilities than previous baselines, even in assembly tasks without semantic information.

Ablation Study

Co-creation Space Analysis Co-creation space is a collaborative environment for all assemblers. In this context,

the parameter k represents the memory capacity of co-creation. This means that k determines how much information each assembler can hold and process. Table 3 displays five different values of k , which range from low to high memory capacities. As the value of k increases, assemblers can hold and process more information. However, there is a trade-off between memory capacity and communication efficiency among assemblers. In this particular task, the optimal result is achieved when $k=10$. This means assemblers with a memory capacity of 10 can hold and process enough information to contribute effectively to the assembly process without becoming overwhelmed. The visual assembly effect, as shown in Figure B.1, demonstrates the impact of memory capacity on the assembly process. At $k=10$, the assembly is best because co-creators can hold and process enough information to assemble the product correctly. However, when $k = x$ (where x is a lower or higher value than 10), problems such as misplaced parts or interference between similar parts can arise due to the information bottleneck. It is worth noting that our method failed to achieve a satisfactory result without co-creat space. Thus, it can be observed that our proposed co-creation space brought outstanding performance improvement consistent with qualitative analysis.

Collision Loss We carried out experiments to assess the effectiveness of the collision loss. As illustrated in Figure 5, we compared the performance of our method with and without the collision loss. The results show a significant enhancement in the model’s capability to place similar parts in different locations with collision loss. This improvement is particularly noticeable when there are numerous comparable parts or a large number of parts. Thus, collision loss is effective to resolving ambiguity between similar parts during the fracture assembly procedure. Furthermore, we add an ablation study for w_c and C on Table 2. The values of w_t , w_r and w_s follow previous works (Sellán et al. 2022).

Coarse-to-fine Coarse-to-fine (CTF) is a sequential process that involves utilizing x different networks to combine distinct parts into a complete shape (Zhan et al. 2020), with each network focusing on more detailed information ($x = 5$ in this paper). This approach enhances efficiency and accuracy by progressively refining the output at each step. The

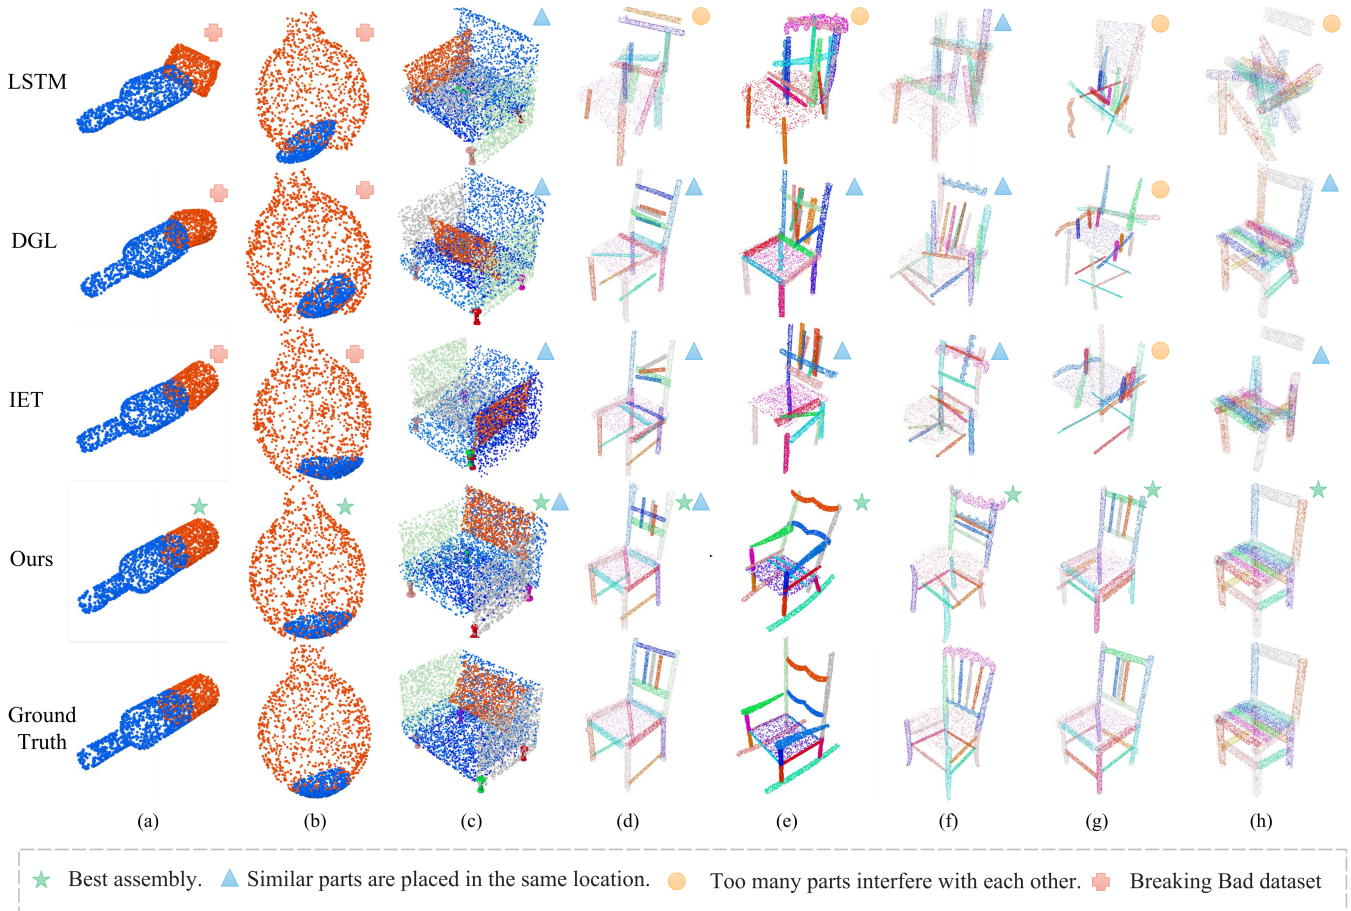


Figure 4: Visual comparisons between our algorithm and the baseline methods on Breaking Bad (a, b) and PartNet dataset (c, d, e, f, g and h).




Figure 5: Collision loss ablation study visualization.

strategy was employed in both the DGL and IET. To verify its impact on experimental results, we implemented a version of the IET that employs CTF with five iterations. As depicted in Table 4, our method still outperformed the others under this condition.

	PA \uparrow	SCD ($\times 10^{-3}$) \downarrow	CA \uparrow
IET	37.50	09.40	24.47
Ours	53.59	07.00	38.97
IET + CTF	41.71	07.80	31.13
Ours + CTF	55.92	06.20	42.69

Table 4: Coarse-to-fine ablation study.

Conclusion

In this paper, we introduce a novel assembly paradigm that effectively addresses scalability issue without relying on semantic information. Firstly, we propose a co-creation space where assemblers compete for write access. Which facilitates step-by-step assembly, reducing confusion when dealing with multiple fractures. Secondly, we have developed a geometric-based collision loss to minimize collision issues during assembly. Our method has been validated across all public datasets, and we are currently exploring its application in robotic hands for assembling real objects.

Acknowledgments

This work is supported by the National Key Research and Development Project of China (2021ZD0110505), National Natural Science Foundation of China (U19B2042), the Zhejiang Provincial Key Research and Development Project (2023C01043), University Synergy Innovation Program of Anhui Province (GXXT-2021-004), Academy Of Social Governance Zhejiang University, Fundamental Research Funds for the Central Universities (226-2022-00064). Furthermore, we thank the anonymous reviewers for their valuable comments and suggestions.

References

- Achlioptas, P.; Diamanti, O.; Mitliagkas, I.; and Guibas, L. 2018. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, 40–49. PMLR.
- Ahuja, K.; Caballero, E.; Zhang, D.; Gagnon-Audet, J.-C.; Bengio, Y.; Mitliagkas, I.; and Rish, I. 2021. Invariance principle meets information bottleneck for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34: 3438–3450.
- Baars, B. J. 1993. *A cognitive theory of consciousness*. Cambridge University Press.
- Chen, Y.-C.; Li, H.; Turpin, D.; Jacobson, A.; and Garg, A. 2022. Neural shape mating: Self-supervised object assembly with adversarial shape priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12724–12733.
- Dehaene, S.; Lau, H.; and Kouider, S. 2021. What is consciousness, and could machines have it? *Robotics, AI, and Humanity: Science, Ethics, and Policy*, 43–56.
- Funkhouser, T.; Shin, H.; Toler-Franklin, C.; Castañeda, A. G.; Brown, B.; Dobkin, D.; Rusinkiewicz, S.; and Weyrich, T. 2011. Learning how to match fresco fragments. *Journal on Computing and Cultural Heritage (JOCCH)*, 4(2): 1–13.
- Goyal, A.; Didolkar, A.; Lamb, A.; Badola, K.; Ke, N. R.; Rahaman, N.; Binas, J.; Blundell, C.; Mozer, M.; and Bengio, Y. 2021. Coordination among neural modules through a shared global workspace. *arXiv preprint arXiv:2103.01197*.
- Grason, G. M. 2016. Perspective: Geometrically frustrated assemblies. *The Journal of Chemical Physics*, 145(11): 110901.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jones, B.; Hildreth, D.; Chen, D.; Baran, I.; Kim, V. G.; and Schulz, A. 2021. Automate: A dataset and learning approach for automatic mating of cad assemblies. *ACM Transactions on Graphics (TOG)*, 40(6): 1–18.
- Ke, N. R.; ALIAS PARTH GOYAL, A. G.; Bilaniuk, O.; Binas, J.; Mozer, M. C.; Pal, C.; and Bengio, Y. 2018. Sparse attentive backtracking: Temporal credit assignment through reminding. *Advances in neural information processing systems*, 31.
- Kirsch, L.; Kunze, J.; and Barber, D. 2018. Modular networks: Learning to decompose neural computation. *Advances in neural information processing systems*, 31.
- Lee, J.; Kim, J.; Chung, H.; Park, J.; and Cho, M. 2020. Fragment Relation Networks for Geometric Shape Assembly. In *Learning Meets Combinatorial Algorithms at NeurIPS2020*.
- Lee, J.; Kim, J.; Chung, H.; Park, J.; and Cho, M. 2022. Learning to Assemble Geometric Shapes. *arXiv preprint arXiv:2205.11809*.
- Lee, Y.; Hu, E. S.; and Lim, J. J. 2021. IKEA furniture assembly environment for long-horizon complex manipulation tasks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 6343–6349. IEEE.
- Li, Y.; Mo, K.; Shao, L.; Sung, M.; and Guibas, L. 2020. 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*, 664–682. Springer.
- Liu, D.; Shah, V.; Boussif, O.; Meo, C.; Goyal, A.; Shu, T.; Mozer, M.; Heess, N.; and Bengio, Y. 2022. Stateful active facilitator: Coordination and Environmental Heterogeneity in Cooperative Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2210.03022*.
- Liu, J.; Hao, J.; Lin, H.; Pan, W.; Yang, J.; Feng, Y.; Wang, G.; Li, J.; Jin, Z.; Zhao, Z.; et al. 2023a. Deep learning-enabled 3D multimodal fusion of cone-beam CT and intraoral mesh scans for clinically applicable tooth-bone reconstruction. *Patterns*, 4(9).
- Liu, J.; Hu, T.; Zhang, Y.; Feng, Y.; Hao, J.; Lv, J.; and Liu, Z. 2023b. Parameter-Efficient Transfer Learning for Medical Visual Question Answering. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- Mo, K.; Zhu, S.; Chang, A. X.; Yi, L.; Tripathi, S.; Guibas, L. J.; and Su, H. 2019. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 909–918.
- Narayan, A.; Nagar, R.; and Raman, S. 2022. RGL-NET: A recurrent graph learning framework for progressive part assembly. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 78–87.
- Pfeiffer, J.; Ruder, S.; Vulić, I.; and Ponti, E. M. 2023. Modular deep learning. *arXiv preprint arXiv:2302.11529*.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2008. The graph neural network model. *IEEE transactions on neural networks*, 20(1): 61–80.
- Schor, N.; Katzir, O.; Zhang, H.; and Cohen-Or, D. 2019. Componet: Learning to generate the unseen by part synthesis and composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8759–8768.
- Sellán, S.; Chen, Y.-C.; Wu, Z.; Garg, A.; and Jacobson, A. 2022. Breaking Bad: A Dataset for Geometric Fracture and

- Reassembly. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Sellán, S.; Luong, J.; Da Silva, L. M.; Ramakrishnan, A.; Yang, Y.; and Jacobson, A. 2021. Breaking Good: Fracture Modes for Realtime Destruction. *ACM Transactions on Graphics (TOG)*.
- Sung, M.; Su, H.; Kim, V. G.; Chaudhuri, S.; and Guibas, L. 2017. ComplementMe: Weakly-supervised component suggestions for 3D modeling. *ACM Transactions on Graphics (TOG)*, 36(6): 1–12.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, 1–5. IEEE.
- Toler-Franklin, C.; Brown, B.; Weyrich, T.; Funkhouser, T.; and Rusinkiewicz, S. 2010. Multi-feature matching of fresco fragments. *ACM Transactions on Graphics (TOG)*, 29(6): 1–12.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wu, R.; Zhuang, Y.; Xu, K.; Zhang, H.; and Chen, B. 2020a. Pq-net: A generative part seq2seq network for 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 829–838.
- Wu, T.; Ren, H.; Li, P.; and Leskovec, J. 2020b. Graph information bottleneck. *Advances in Neural Information Processing Systems*, 33: 20437–20448.
- Zhan, G.; Fan, Q.; Mo, K.; Shao, L.; Chen, B.; Guibas, L. J.; Dong, H.; et al. 2020. Generative 3d part assembly via dynamic graph learning. *Advances in Neural Information Processing Systems*, 33: 6315–6326.
- Zhang, R.; Kong, T.; Wang, W.; Han, X.; and You, M. 2022. 3D Part Assembly Generation With Instance Encoded Transformer. *IEEE Robotics and Automation Letters*, 7(4): 9051–9058.
- Zhu, D.; Li, Y.; Yuan, J.; Li, Z.; Kuang, K.; and Wu, C. 2023a. Universal domain adaptation via compressive attention matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6974–6985.
- Zhu, D.; Li, Y.; Zhang, M.; Yuan, J.; Liu, J.; Kuang, K.; and Wu, C. 2023b. Bridging the Gap: Neural Collapse Inspired Prompt Tuning for Generalization under Class Imbalance. *arXiv preprint arXiv:2306.15955*.