

This ICCV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Towards realistic symmetry-based completion of previously unseen point clouds

Taras Rumezhak^{1,2} Oles Dobosevych¹ Rostyslav Hryniv¹ Vladyslav Selotkin² Volodymyr Karpiv² Mykola Maksymenko²

¹The Machine Learning Lab at Ukrainian Catholic University ²SoftServe Inc.

{rumezhak, dobosevych, rhryniv}@ucu.edu.ua

{trume, vselo, vkarpi, mmaks}@softserveinc.com

Abstract

3D scanning is a complex multistage process that generates a point cloud of an object typically containing damaged parts due to occlusions, reflections, shadows, scanner motion, specific properties of the object surface, imperfect reconstruction algorithms, etc. Point cloud completion is specifically designed to fill in the missing parts of the object and obtain its high-quality 3D representation. The existing completion approaches perform well on the academic datasets with a predefined set of object classes and very specific types of defects; however, their performance drops significantly in the real-world settings and degrades even further on previously unseen object classes.

We propose a novel framework that performs well on symmetric objects, which are ubiquitous in man-made environments. Unlike learning-based approaches, the proposed framework does not require training data and is capable of completing non-critical damages occurring in customer 3D scanning process using e.g. Kinect, time-of-flight, or structured light scanners. With thorough experiments, we demonstrate that the proposed framework achieves stateof-the-art efficiency in point cloud completion of real-world customer scans. We benchmark the framework performance on two types of datasets: properly augmented existing academic dataset and the actual 3D scans of various objects. The code is available here: https://github.com/softserveincrnd/symmetry-3d-completion

1. Introduction

The ability to recognize the geometry of a 3D object is a key pre-requisite in many computer vision applications. This task can be accomplished with 3D scanners, including LIDARs, time-of-flight and structured light sensors that have seen tremendous progress over the last few years and have become more accessible and easy to use.

Point cloud format is commonly used for the 3D shape representation of an object. Due to the limited sensor reso-



Figure 1. Method performance pipeline: a) damaged object, b) completed regions (green) c) result, d) original object

lution and occlusions, an incomplete point cloud is typically generated as an output of a scanning pipeline. Recovering a damaged point cloud by filling in the missing parts, called *point cloud completion*, is the step required in many 3D scanning pipelines as it generates a more accurate 3D representation of the object and moreover enhances the overall performance of further 3D computer vision pipelines.

Most existing solutions of the point cloud completion task are based on deep neural networks trained on large datasets with a limited number of classes that the solution can complete. State-of-the-art approaches use datasets with e.g. up to 40 classes (ShapeNet [3]), which is a serious constraint for real-world applications as reconstructing new types of objects requires a long training procedure and massive datasets that are hard to collect.

To overcome this problem, we propose a novel framework that does not require training data and is capable of completing previously unseen symmetric objects. The symmetry of interest here is a mirror reflection symmetry that is present in most human artefacts.

Our use case is different from those related e.g. to autonomous driving and for which some solutions have been proposed since they are mostly concentrated on completion of partially observed objects as in the KITTI dataset [11]. Our framework concentrates on a different use case with imperfect full object scans made with time-of-flight or structured light scanners, in which the damages can occur when the speed of scanner motion is too high or due to the shadows and occlusions. For example, with the True Depth camera, some regions can be missed when the speed of iPhone motion is greater than the point perception speed or when the object is on a turntable and images are taken by one camera. Another example is the structured light scanning pipeline where the projector as the source of light is set on the position appearing on one side of the scanning object and resulting shadows can lead to the damages in the resulting point cloud as seen in Figure 1.

Our framework consists of three main steps. On the *first* step, we approximate the symmetry plane of the given point cloud P using the Principal Component Analysis (PCA) [9] applied to the surface normals in combination with a novel approach of Dominant convex hull directions. Then we construct the mirror reflection P' of P with respect to the approximated symmetry plane. Due to the broken symmetry in P and imperfect symmetry plane approximation, P' does not coincide with P. The second step consists in the registration of the point clouds P and P' that would minimize the mismatch. We apply the Global registration algorithm of feature point-to-point matching via Fast Point Feature Histogram (FPFH) [17] along with Random Sample Consensus (RANSAC) with parameters suggested by Choi et al. [4]. The registration is iteratively improved with the Iterative Closest Point (ICP) algorithm [2]. On the third step, we detect the missing parts of the point cloud P from its optimal mirror reflection P' using the proposed proximity measure and finish the completion of the point cloud P to better represent the initial 3D object.

Performance of the proposed framework is demonstrated on point cloud completion of two datasets, including the standard academic dataset ShapeNet which was properly augmented for our use case with non-critical damage-rate as well as 3D scans of real-world objects collected for the purpose of this research. The accuracy of our algorithm was compared with several neural networks, such as GR-Net: Gridding Residual Network (GRNet) [26], Morphing and Sampling Network (MSN) [13], Point Completion Network (PCN) [28]. In addition to not requiring any pretraining and being not restricted to particular object classes, our approach also achieves state-of-the-art completion results comparable to those obtained by other methods. Since the symmetry plane detection accuracy is critical for the overall performance of the framework, we also benchmark it on the 2017 ICCV Challenge: Detecting Symmetry in the Wild [10] dataset and report better or comparable results with the solution of Cicconet et al. [5], which was the only one with the provided code. Finally, we demonstrate that our framework generalizes to previously unseen real-world objects collected by three types of 3D scanners: Microsoft Kinect, iPhone's TrueDepth camera (or LIDAR in the latest versions) and the structured light scanner.

The main contributions of the paper can be summarized as follows:

- a learning-free framework is proposed for point cloud completion based on the mirror symmetry;
- a novel symmetry plane detection technique is introduced as an intermediate step of the framework;
- a novel distance metric is developed to measure the accuracy of point cloud matching;
- the experiments were conducted on the common academic dataset ShapeNet and on the newly collected real-world dataset of scans; the benchmarks indicate that the proposed framework performs well on the general class of symmetric objects and even outperforms state-of-the-art approaches on the provided datasets.

2. Related work

2.1. Point cloud completion

Learning-based. The pioneering approaches to point cloud completion are mainly focused on the usage of deep neural networks (NN). Several approaches are based on the multi-layer perceptron (MLP) due to its simplicity and solid representation power. Qi et al. [16] introduce PointNet that is applied directly to a point cloud and correctly handles the permutation invariance of its input. At the same time, this class of methods aggregates features with symmetric functions, such as Max-pooling, which do not fully exploit the geometric structure of a point cloud. Xie et al. [26] solve the problem of the structural and contextual information loss by introducing 3D grids as intermediate representations as a regularization to unordered point clouds and based on such grids present the Gridding Residual Network (GRNet) which shows great results on the Completion3D benchmark [23]. Also to overcome the structure information loss Groueix et al. [12] introduced the method that represents a 3D shape as a collection of parametric surface elements and utilizes surface representation of the shape instead of voxel grids or point clouds. Similarly, Liu et al. [13] propose the method that also uses a collection of parametric surface elements. In the first stage they predict a complete coarse-grained point cloud and in the second stage the algorithm merges the predictions with the input point cloud with the proposed sampling method. In contrast to intermediate parametrization or voxelization, Yuan et al. [28] propose a Point Completion Network (PCN) which operates directly on the point clouds without any structural or geometry assumption and completes them in a coarse-to-fine fashion.

Geometry-based. Besides learning-based approaches some attempts have been made with geometry-based methods which complete the shapes using geometry features without any external data. The geometry-based methods themselves can be categorized into surface interpolations and symmetry-driven methods. Sarkar et al. [18] propose a 3D shape parameterization method by surface patches that performs interpolated shape completion of complicated surface textures. The results achieved by Sorkine et al. [20] used the connectivity of the mesh containing the geometric information that approximates a set of the control point in a least-squares manner and Zhao et al. [29] propose the method that can fill great holes in triangular mesh models by approximation of newly created normals. We should note that the primary task of these papers is to fill small locally uncompleted regions by smooth interpolation and the information based on a mesh, which is not the use case with point clouds. Mitra et al. [14] discussed the strengths and limitations of existing 3D symmetry algorithms. Pauly et al. [15] introduced the framework for detecting the repeated geometric patterns in mesh-based models by the analysis of pairwise similarity transformations. Sung et al. [22] method uses a 3D model collection to build the structural part-based priors to perform the completion. They also focus on the use case with low-quality consumer-level scanners and achieve compatible results, but still are the training-based approach.

2.2. Symmetry plane detection

Application of the symmetry plane detection for point clouds can be found in various fields of computer vision and computer graphics. Fan Xue et al. [27] showed that symmetry estimation can be well accomplished with prior knowledge of the object type. However, in the general case, the object type is unknown beforehand. With an assumption of a bilateral object (having perfect or imperfect mirror symmetry) Combes et al. [7] presented an iterative approach using the maximum likelihood principle and the expectation-maximization (EM) algorithm. To deal with the problem of noisy or missing data, Sipiran et al. [19] developed a feature-based approach. The algorithm generates a set of symmetry candidates based on detected features. The vote-based system validates all candidates and selects the best one. Several interesting methods were presented on the 2017 ICCV Challenge [10]: Detecting Symmetry in the Wild. Cicconet et al. [5] developed a simple and efficient symmetry estimation approach based on registration which is also a key step in our pipeline. However, the accuracy of the algorithm drops when the points in the symmetric parts

are not perfectly aligned or missing parts are presented in the point clouds. Speciale et al. [21] propose the approach of detecting symmetries from incomplete data which then allows the scene completion. They extend standard symmetry detection techniques to exploit partially unexplored domains. Cohen et al. [6] propose a method to recover symmetry relations using the geometry cues. The symmetry priors are incorporated in a new constrained bundle adjustment formulation. The method can perform 3D completion to get rid of Structure-from-Motion (SfM) artifacts based on the symmetries.

3. Pipeline description

Given a real-world 3D object D that is symmetric with respect to a plane $\pi \subset \mathbb{R}^3$. The symmetry plane $\pi = \pi(\mathbf{p}, \mathbf{n})$ is uniquely determined by any of its points $\mathbf{p} \in \pi$ and a unit normal \mathbf{n} ; then the mirror reflection T_{π} with respect to π is given by $T_{\pi}\mathbf{x} = \mathbf{x} - 2\langle \mathbf{x} - \mathbf{p}, \mathbf{n} \rangle \mathbf{n}$, where $\langle \cdot, \cdot \rangle$ is the standard scalar product in \mathbb{R}^3 . The object D is called *symmetric* with respect to π if $T_{\pi}D = D$.

In real-world applications, D is represented by a point cloud $P = {\mathbf{x}_1, \ldots, \mathbf{x}_n}$ of points $\mathbf{x}_i \in D$. In the ideal situation, the points of P are uniformly spread inside D so that the mirrored cloud $T_{\pi}P$ is almost identical to P, in the sense that the Chamfer distance between P and $T_{\pi}P$ (5) is small enough. However, in the realistic cases point clouds P representing D have holes, i.e., there are sub-regions $D_j \subset D$, $j = 1, 2, \ldots, h$ with no or very few points \mathbf{x}_i . Our goal is to fill in these holes, i.e., to find a point cloud $P^* \supset P$ that is symmetric with respect to a plane π^* and has the smallest possible volume. In order that P^* would represent the whole D, the missing part $\bigcup_{j=1}^h D_j$ should have a void (or negligibly small) intersection with its mirror image $T_{\pi}(\bigcup_{i=1}^h D_j)$.

The completion framework proposed in this paper (see Fig. 2) includes three main steps: symmetry plane estimation, its iterative approximation, and then localization of the holes and their filling. The first two steps are based on point cloud registration and include several preliminary steps, such as bounding box and centroid computation, symmetry plane candidates computation and approximation.

3.1. Bounding box and centroid computation

Missing regions in the point cloud induce errors to the whole objects geometry estimation. To localize the point cloud and minimize the influence of its missing parts, we put it inside its *bounding box*; this is a cuboid of minimal volume that is circumscribed around the point cloud and whose faces are parallel to the coordinate planes. It is constructed by detecting the extreme points of the point cloud along the coordinate x, y, and z-axes; the coordinates of the cuboid vertices are combinations of these extreme values in



Figure 2. Pipeline of the algorithm

each direction:

$$\mathbf{v}_1 = (x_{\min}, y_{\min}, z_{\min}),$$

$$\mathbf{v}_2 = (x_{\max}, y_{\min}, z_{\min}),$$

$$\dots$$

$$\mathbf{v}_8 = (x_{\max}, y_{\max}, z_{\max}).$$

The *centroid* **c** of the bounding box has coordinates

$$\mathbf{c} = \left(\frac{x_{\max} + x_{\min}}{2}, \frac{y_{\max} + y_{\min}}{2}, \frac{z_{\max} + z_{\min}}{2}\right).$$
 (1)

The bounding box centroid **c** is more stable to point cloud damages than the point cloud mass center mass center $\overline{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_i$ and we prefer it as a reference point for the symmetry plane approximation in the dominant convex hull direction method of Subsection 3.2.3.

3.2. Symmetry plane estimation

Standard approaches to symmetry plane estimation fail to find the correct symmetry plane π due to random point sampling and non-symmetric missing parts in the point cloud. However, the proposed algorithm requires only an approximation of the symmetry plane that is close enough to the correct one. We suggest two approaches to the symmetry plane estimation, which are based on the principal component analysis (PCA) applied to the surface normals and dominat convex hull directions, respectively. These two methods generate several candidates for the symmetry plane, and each candidate is evaluated with the balance distance metrics (3) to select the best one. As experiments show, using this best candidate as initial symmetry plane approximation results in solutions with better Chamfer distance to the ground truth (see Fig. 3).



Figure 3. Comparison of the final results with a random initial symmetry plane approximation and suggested approximation using PCA candidates. The experiments were run on the properly augmented ShapeNet dataset with 9 different damage rates

3.2.1 Principal component analysis

Principal component analysis (PCA) can be used to estimate the symmetry plane of a point cloud P. The PCA calculates the eigenvalues and eigenvectors (principal components) of the 3×3 covariance (i.e., inertia) matrix Σ of P. If π is a symmetry plane of P, then the principal components c are invariant under T_{π} , so that each principal component either belongs to π or is orthogonal to it.

Therefore, given a point cloud P that is not completely symmetric with respect to π , we suggest a candidate π^* for π as a plane through the bounding box centroid c of (1) that is orthogonal to a principal component. Assuming that Σ has no repeated eigenvalues (which is the case when the point cloud has no multiple planes of symmetry), we thus get three candidates π^* of which the best one is chosen.

Experiments showed that the PCA is not sufficiently robust to random damages; therefore, we apply it not to the initial point cloud P itself but to two point clouds that are projections onto the unit sphere S^2 of the surface normals and the so called dominant convex hull directions discussed in Subsections 3.2.2 and 3.2.3.

3.2.2 PCA of normal directions

We use the open3d.geometry.estimate_normals algorithm of Open3D² library [31] to calculate the point

²http://www.open3d.org/

cloud normals \mathbf{n}_j by locally fitting planes through the surface points of P and register them on the unit sphere S^2 . Then we run the PCA on the resulting set N of normals to find a candidate π^* of the symmetry plane π as explained in Subsection 3.2.1.

In the case where the point cloud P was obtained from the uniformly representing point cloud of a symmetric object D by removing several randomly located holes, the PCA of normal directions in most cases returns a good approximation π^* for the ground-truth symmetry plane π .

3.2.3 Dominant convex hull directions

Our experiments demonstrated that there are cases when the PCA of normal directions fails to suggest a good candidate for the symmetry plane π . A reason for that is that the missing parts D_j of the point cloud P may significantly affect the surface and thus the normals. To overcome this problem, we propose a complementary approach for symmetry plane estimation using the PCA of the dominant directions in the point cloud convex hull.

Firstly, the convex hull is built around the point cloud P using the Quickhull algorithm [1] implemented in Open3D library. It consists of numerous intervals joining points on the point cloud surface, and we register their directions $\pm \mathbf{u}_j$ on the unit sphere S^2 . If the point cloud P has mirror symmetry π , then the set of unit directions \mathbf{u}_j will be almost symmetric with respect to π .

In presence of measurement errors and missing parts D_j , the point cloud P will not be symmetric with respect to π but the set U of unit directions on the sphere S^2 is typically close to symmetric, see Figure 4. We apply the PCA to the set U of directions on the sphere S^2 and get three candidates for a symmetry plane as described in Subsection 3.2.1.



Figure 4. a) Dominant convex hull directions: Projections of the convex hull lines onto the unit sphere S^2 ; b) Projections of the points normals onto the unit sphere S^2

3.3. Metric

To measure the symmetry plane estimation accuracy and to choose the best among several candidates, a new metric is suggested. For a point cloud P and a symmetry plane candidate σ , we first calculate the mirror symmetric image $P' = T_{\sigma}(P)$ of P. Denote by $Q(\mathbf{x}, d)$ the cube of center \mathbf{x} and side length d, whose faces are parallel to the coordinate planes; we call a point $\mathbf{x} \in P \cup P'$ balanced if

$$\frac{\left|\left|Q(\mathbf{x},d)\cap P\right| - \left|Q(\mathbf{x},d)\cap P'\right|\right|}{\left|Q(\mathbf{x},d)\cap P\right| + \left|Q(\mathbf{x},d)\cap P'\right|} \le \varepsilon$$
(2)

for experimentally established threshold ε ; here |A| denotes the cardinality of a set A. Denote by B(P, P') the number of all balanced points in $P \cup P'$; then the *balanced distance* BD(P, P') between the point clouds P and P' is

$$BD(P,P') = 1 - \frac{B(P,P')}{|P| + |P'|};$$
(3)

this is an analogue of the standard *intersection-over-union* measure.

3.4. Registration

The registration of the point cloud is required since the candidate for the symmetry plane is only an estimation. The common method for point cloud alignment is Iterative Closest Point (ICP), which is known as the local registration method because it relies on a rough initial alignment. In contrast, we will use global registration which does not require the initial alignment but produces a less accurate result, so it will be used as the input of ICP which will perform the post-processing refinement.

3.4.1 Normals and FPFH features

The FPFH features [17] describe the local geometric properties of a point cloud and are constructed of the vectors with 33-dimensions. Computation of the FPFH features requires point normal estimation, which is based on a local neighborhood of the point and involves covariance analysis. The camera position or the alignment axis should be provided to choose one of the two directions consistently.

3.4.2 Global registration

Extracted FPFH features can be used to perform the registration of the point clouds. The registration based on the RANSAC algorithm with parameters proposed by Choi et al. [4] is slow due to numerous model proposals and evaluations, so it is performed on the down-sampled point cloud. The Fast Global Registration [30] algorithm optimizes the process of finding point correspondences as for each iteration it has no model proposal and evaluation involved.

3.4.3 Refine registration via ICP

Iterative Closest Point [2] is a classical approach for point cloud registration. The algorithm finds the best rotation matrix R and translation vector t to minimize the squared error defined as

$$E(R,t) = \frac{1}{N_p} \sum_{i=1}^{N_p} ||\mathbf{x}_i - R\mathbf{p}_i - \mathbf{t}||^2, \qquad (4)$$

where \mathbf{x}_i and \mathbf{p}_i are the corresponding points in point clouds P and P' to be aligned, and N_p is the number of corresponding point pairs.

If the correct correspondences are known, then the alignment can be performed accurately in one step. However, usually, the matching points are not known beforehand, and the ICP finds the optimal transformations iteratively.

3.5. Hole detection and filling

After the best point cloud alignment of P and its mirrored image P' has been achieved (i.e., the best symmetry plane π^* proposed), the hole detection algorithm is applied. Every point \mathbf{x}' in the mirrored point cloud $P' = T_{\pi^*}(P)$ that is not balanced in the sense of (2) is in a hole of D if $|Q(\mathbf{x}', d) \cap P'|$ is bigger than $|Q(\mathbf{x}', d) \cap P|$. The holes in the original point cloud are detected this way along with the corresponding parts of the mirrored point cloud P' that should be added to the point cloud P to finish completion.

3.6. Skipping

Performance analysis of the proposed pipeline showed that there are some cases where the completed point clouds have significantly bigger Chamfer distance (5) to the ground truth than the original one. One of the main reasons for this is that not all objects in the datasets have mirror symmetry; in some other cases the algorithm fails to suggest a good symmetry plane approximation.

To deal with this issue, we add a validation step that decides whether the completed point cloud P^* or the original damaged point cloud P should be returned. After the completion, we calculate the scaled Chamfer distance $CD(P, P^*)/s$ between P and P^* ; if it exceeds a threshold value d^* , then we expect that the Chamfer distance between the completion P^* and the ground truth will be large; we thus regard P^* as a poor reconstruction and return the original P instead. This step lets us judge whether the object is symmetric with non-critical damages.

The threshold value d^* was chosen experimentally based on the corresponding performance curve of d^* . The scale *s* is chosen as the average distance between the closest points in the input point cloud *P*.



Figure 5. Comparison of methods on damaged ShapeNet objects.



Figure 6. Comparison of methods on synthetic data with: a) 15%, b) 35% damage rate, and real scans with classes on which compared approaches were: c) trained on, d) not trained on.

4. Experiments

For the evaluation purpose, the ShapeNet dataset [3] was chosen as it is the most frequently used source for 3D completion applications. We first explain the way how evaluation dataset was formed, then describe the metrics used for performance evaluation of the proposed algorithm and comparison with other methods on point cloud completion and mirror symmetry detection. We then report the corresponding scores for both tasks and, finally, show point cloud completion performance of the proposed method on real-world 3D object scans in the wild.

4.1. Datasets creation

Performance evaluation is measured on an appropriately augmented ShapeNet dataset which is composed of 8 categories with 30 974 3D models already uniformly sampled with 16 384 points on the mesh. In some popular benchmarks (e.g., Completion3D [25]) the mentioned dataset was



Figure 7. Comparison of methods on the ShapeNet based dataset with different damage rates.

Method	5%	15%	25%	35%	45%
MSN	11.32	11.16	11.23	11.37	11.42
PCN	10.63	10.60	10.61	10.57	10.70
GRNet	8.95	8.83	8.98	9.16	9.12
Ours (-Skip)	3.27	8.43	15.43	19.42	25.67
Ours	2.92	5.44	7.56	10.59	13.78

Table 1. Numerical Chamfer Distance comparison based on different damage rate percentage. (-Skip) - without Skipping step.

damaged sufficiently with a damaged rate (a ratio of a missing surface) up to 70%, which is a normal case in autonomous driving when LIDAR captures only one side of the car and the algorithm must complete the whole unseen region. As our case is different, the type of damages must also differ from the existing benchmark. We expect to have the objects with 15-30% damage rate, which is natural for our case. To perform experiments of various difficulty, we created nine versions of dataset with the damage rate DRfrom 5% to 45% with the step of 5% and a random number of missing regions fluctuating between |0.7DR| and [0.95DR]. The random number of damaged regions, their size and location is a great test for the robustness of our algorithm. In learning-based approaches, robustness could be achieved by noise-based data augmentation, while in our case random damaged regions are the inverse of the noise.

4.1.1 Real-world data

To assess framework performance on real data, we also collected a dataset of more than 200 scanned objects. Large objects such as a chair were scanned with the Microsoft Kinect v2, while small objects such as a toy plane were scanned with the iPhone's TrueDepth camera (or LIDAR) or our structured light 3D scanner.

Learning-based approaches were trained on the academic dataset and do not generalize well to the real data. To fit the input standard of compared approaches, the obtained objects were downsampled, properly scaled and translated accordingly to contain the fixed number of points similar to the original ShapeNet dataset. Even on the categories the compared approaches were trained on, their performance remains poor. While our framework is not limited to a given type and size of point clouds or a defect so it successfully completes 3D scans of real objects as shown in the Fig. 6.

4.1.2 Symmetry Dataset

One of the key steps of our framework is the estimation of the best symmetry plane. To evaluate the performance of the proposed algorithm on this step, we benchmarked it on the 3D Global Reflection Symmetry dataset from 2017 ICCV Challenge: Detecting Symmetry in the Wild [10] consisting of 1 795 models with up to three symmetry planes in each.

4.2. Evaluation metrics

4.2.1 Chamfer distance

As a proximity measure between two point clouds P_1 and P_2 , we use the Chamfer distance [8, 24] that calculates the average closest point distance:

$$CD(P_1, P_2) = \frac{1}{|P_1|} \sum_{\mathbf{x} \in P_1} \operatorname{dist}(\mathbf{x}, P_2) + \frac{1}{|P_2|} \sum_{\mathbf{y} \in P_2} \operatorname{dist}(\mathbf{y}, P_1)$$
(5)

4.2.2 Symmetry Plane Detection Accuracy

As proposed in the ICCV 17 Challenge [10], the 3D planar reflective symmetries are evaluated according to the position and orientation of the symmetry plane with respect to the ground truth. The detected symmetry is considered correct if both the position of the reference point is on the plane and the normal is close to ground truth with some threshold. The symmetry is rejected if the cosine distance between the symmetry normals is above a threshold: $\arccos(|n \cdot n_{\rm GT}|) > \theta$. A symmetry is also rejected if the distance of the tested center c to the ground truth plane restricted to the bounding parallelogram is larger than a threshold $\|\mathbf{c} - \Pi_{\rm GT}(\mathbf{c})\|_2 > \tau$. Accuracy is the relative frequency of correct predictions among all objects.

4.3. Comparison with other methods

4.3.1 Point Cloud Completion

We evaluate the completion results on the same datasets subject to different damage rates. GRNet, the top method of the Completion3D benchmark, shows good results in general, but approximately 17% of its outputs have a very large Chamfer distance thus leading to a moderate average metric. MSN gives worse results than GRNet and still shows

Method	Original	Damaged by 10%	Damaged by 20%	Damaged by 30%	Damaged by 40%
Symmetry via Registration	68.63%	45.75%	44.44%	42.44%	35.29%
Ours	66.24%	50.32%	49.67%	48.36%	45.09%

Table 2. Accuracy comparison with Symmetry via Registration with threshold $\theta = 0.2$

worse performance on small damages than our method. Our framework has an 8% of anomaly big metric values due to a poor symmetry plane detection or its absence in the original object; applying the Skipping validation, we significantly increase the average results.

Weak sensitivity of GRNet, PCN and MSN to damage rate can be explained by the fact that the corresponding neural networks are trained on particular classes of the Completion 3D benchmark dataset (based on ShapeNet) and thus can complete point clouds of the learned objects with even 70-80% of damage. Our method, on the contrary, would not be able to successfully reconstruct so heavily damaged object; however, scans of so poor quality are highly unlikely to be produced in a multiview scanning. Even without Skipping validation, our results are comparable with stateof-the-art approaches on objects with at most 20% damage rate, which is quite natural for the real world, while the final result is better with up to 30% damage rate and degrades only after 35 - 40%. PCN and GRNet results are practically the same, see Fig. 7, however, PCN overall accuracy is lower. Also, see the visual comparison on the damaged ShapeNet objects in Fig. 5.

4.3.2 Symmetry Plane Detection Accuracy

We have compared the achieved symmetry plane estimation accuracy to the Symmetry via Registration (SvR) algorithm [5]. SvR is based on the optimal symmetric pairwise assignment of curves. This algorithm outperforms ours on the non-damaged objects by 2.4%. However, its quality significantly deteriorates with damage rate growth, while the proposed method gives more stable results; see Tab. 2.

5. Limitations

There are limitations in our approach: non-symmetric objects, symmetric damages, and their sizes. We need to notice that in the case of trained models non-symmetric objects could be addressed. However, each such non-symmetric class has to be introduced separately, while our method is readily applicable to new symmetric objects. Symmetric damages with a damage rate up to 30% are not likely, but in our case, especially in the structured light pipeline, such damages are highly unlikely as shadows or glares can appear only on one side of the object. For our use case, we focus on the moderate damage rate of 15 to 30% which is normal for multiview structured light or hand SfM

scanners. Another must mention point is the time performance of our method. For the 16K points ShapeNet model it takes 1.3 seconds on the Python CPU implementation to get the result while the compared approaches do not support CPU.

6. Conclusion

Point cloud completion is a critical step in real-world 3D scanning pipelines where objects are typically damaged due to the imperfections of the scanning process or reconstruction algorithms. Being an active field of research, the latest SOTA results are often associated with trained NN-based models that learn peculiarities of different classes of objects. While for specific cases such approaches showcase impressive results in the general case of new types or unseen objects the geometry-based methods can show more robust performance.

In this paper we present a novel point cloud completion framework that exploits symmetries of the objects to fill in the damaged parts of the objects. The method requires no training data and yet provides better or comparable performance than recent SOTA NN-based approaches such as PCN, MSN or GRNet.

First, we explore the robustness of our algorithm on academic benchmark datasets that were used to train the above models. Here for low damage rates below 30% our method clearly outperforms other approaches and showcases Chamfer distance of 3 to 9. For the higher damage rates up to 45%, it shows comparable to SOTA performance. Next, we present a new real-world dataset of more than 200 objects and showcase qualitatively that our approach is robust to variation in object shape, while trained methods in some cases fail on previously unseen data (see Fig. 6). The robustness of our pipeline is linked to a novel method of symmetry plane approximation that we outline and benchmark in the manuscript. In comparison to the benchmark SvR algorithm it provides a significant boost in performance on damaged point clouds as showcased in Table 2. Finally, to reproduce our results, we provide the code of algorithm implementation.

As the proposed Point Cloud completion pipeline was integrated and tested in a production 3D scanning setting, we believe that it has a wide potential as a robust engineering solution for a wide range of practical applications.

References

- C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22(4):469–483, Dec. 1996. 5
- [2] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. 2, 6
- [3] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 1, 6
- [4] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. 06 2015. 2, 5
- [5] M. Cicconet, D. G. C. Hildebrand, and H. Elliott. Finding mirror symmetry via registration and optimal symmetric pairwise assignment of curves. In 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), pages 1749–1758, 2017. 2, 3, 8
- [6] A. Cohen, C. Zach, S. N. Sinha, and M. Pollefeys. Discovering and exploiting 3d symmetries in structure from motion. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, pages 1514–1521, 2012. 3
- [7] Benoît Combès, Robin Hennessy, John Waddington, Neil Roberts, and Sylvain Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. 26th IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 06 2008. 3
- [8] Haoqiang Fan, Hao Su, and Leonidas Guibas. A point set generation network for 3d object reconstruction from a single image, 2016. 7
- [9] Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. 2
- [10] C. Funk, S. Lee, M. R. Oswald, S. Tsogkas, W. Shen, A. Cohen, S. Dickinson, and Y. Liu. 2017 iccv challenge: Detecting symmetry in the wild. In 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), pages 1692–1701, 2017. 2, 3, 7
- [11] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 2
- [12] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. Atlasnet: A papiermâché approach to learning 3d surface generation, 2018. 2
- [13] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion, 2019. 2
- [14] Niloy Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. Symmetry in 3d geometry: Extraction and applications. *Computer Graphics Forum*, 32, 09 2013. 3
- [15] Mark Pauly, Niloy J. Mitra, Johannes Wallner, Helmut Pottmann, and Leonidas J. Guibas. Discovering structural

regularity in 3d geometry. *ACM Trans. Graph.*, 27(3):1–11, Aug. 2008. 3

- [16] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation, 2017. 2
- [17] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In 2009 IEEE International Conference on Robotics and Automation, pages 3212– 3217, 2009. 2, 5
- [18] Kripasindhu Sarkar, Kiran Varanasi, and Didier Stricker. Learning quadrangulated patches for 3d shape parameterization and completion, 2017. 3
- [19] Ivan Sipiran, Robert Gregor, and Tobias Schreck. Approximate symmetry detection in partial 3d meshes. *Comput. Graph. Forum*, 33(7):131–140, Oct. 2014. 3
- [20] O. Sorkine and D. Cohen-Or. Least-squares meshes. In Proceedings Shape Modeling Applications, 2004., pages 191– 199, 2004. 3
- [21] Pablo Speciale, Martin Oswald, Andrea Cohen, and Marc Pollefeys. A symmetry prior for convex variational 3d reconstruction. 10 2016. 3
- [22] Minhyuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas Guibas. Data-driven structural priors for shape completion. ACM Trans. Graph., 34(6), Oct. 2015. 3
- [23] Lyne Tchapmi, Vineet Kosaraju, and Silvio Savarese. Stanford 3d object completion and reconstruction benchmark. *arXiv preprint*, 2019. 2
- [24] L. P. Tchapmi, V. Kosaraju, H. Rezatofighi, I. Reid, and S. Savarese. Topnet: Structural point cloud decoder. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 383–392, 2019. 7
- [25] Lyne P Tchapmi, Vineet Kosaraju, S. Hamid Rezatofighi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6
- [26] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Jiageng Mao, Shengping Zhang, and Wenxiu Sun. Grnet: Gridding residual network for dense point cloud completion, 2020. 2
- [27] Fan Xue, Ke Chen, and Weisheng Lu. Architectural symmetry detection from 3d urban point clouds: A derivative-free optimization (dfo) approach. pages 513–519, 10 2018. 3
- [28] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network, 2019. 2, 3
- [29] Wei Zhao, Shuming Gao, and Hongwei Lin. A robust holefilling algorithm for triangular mesh. volume 23, pages 22– 22, 11 2007. 3
- [30] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. In Computer Vision – ECCV 2016, pages 766–782, 10 2016. 5
- [31] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. arXiv:1801.09847, 2018. 4