
CAD-DEFORM: DEFORMABLE FITTING OF CAD MODELS TO 3D SCANS

Vladislav Ishimtsev*

Skolkovo Institute of Science and Technology
Russia

Alexey Bokhovkin*

Skolkovo Institute of Science and Technology
Russia

Alexey Artemov

Skolkovo Institute of Science and Technology
Russia

Savva Ignatyev

Skolkovo Institute of Science and Technology
Russia

Matthias Niessner

Technical University of Munich, Germany

Denis Zorin

New York University, USA
Skolkovo Institute of Science and Technology
Russia

Evgeny Burnaev

Skolkovo Institute of Science and Technology
Russia

ABSTRACT

Shape retrieval and alignment are a promising avenue towards turning 3D scans into lightweight CAD representations that can be used for content creation such as mobile or AR/VR gaming scenarios. Unfortunately, CAD model retrieval is limited by the availability of models in standard 3D shape collections (*e.g.*, ShapeNet). In this work, we address this shortcoming by introducing CAD-Deform¹, a method which obtains more accurate CAD-to-scan fits by non-rigidly deforming retrieved CAD models. Our key contribution is a new non-rigid deformation model incorporating smooth transformations and preservation of sharp features, that simultaneously achieves very tight fits from CAD models to the 3D scan and maintains the clean, high-quality surface properties of hand-modeled CAD objects. A series of thorough experiments demonstrate that our method achieves significantly tighter scan-to-CAD fits, allowing a more accurate digital replica of the scanned real-world environment while preserving important geometric features present in synthetic CAD environments.

Keywords Scene reconstruction · Mesh deformation

1 Introduction

A wide range of sensors such as the Intel RealSense, Google Tango, or Microsoft Kinect can acquire point cloud data for indoor environments. These data can be subsequently used for reconstructing 3D scenes for augmented and virtual reality, indoor navigation and other applications [1, 2, 3, 4, 5, 6]. However, available 3D reconstruction algorithms are not sufficient for many applied scenarios as the quality of the result may be significantly affected by noise, missing data, and other artifacts such as motion blur found in real scans, disabling reconstruction of fine-scale and sharp geometric features of objects. In most instances, reconstructions are still very distant from the clean, 3D models created manually.

*Equal contribution

¹The code for the project: <https://github.com/alexeybokhovkin/CAD-Deform>

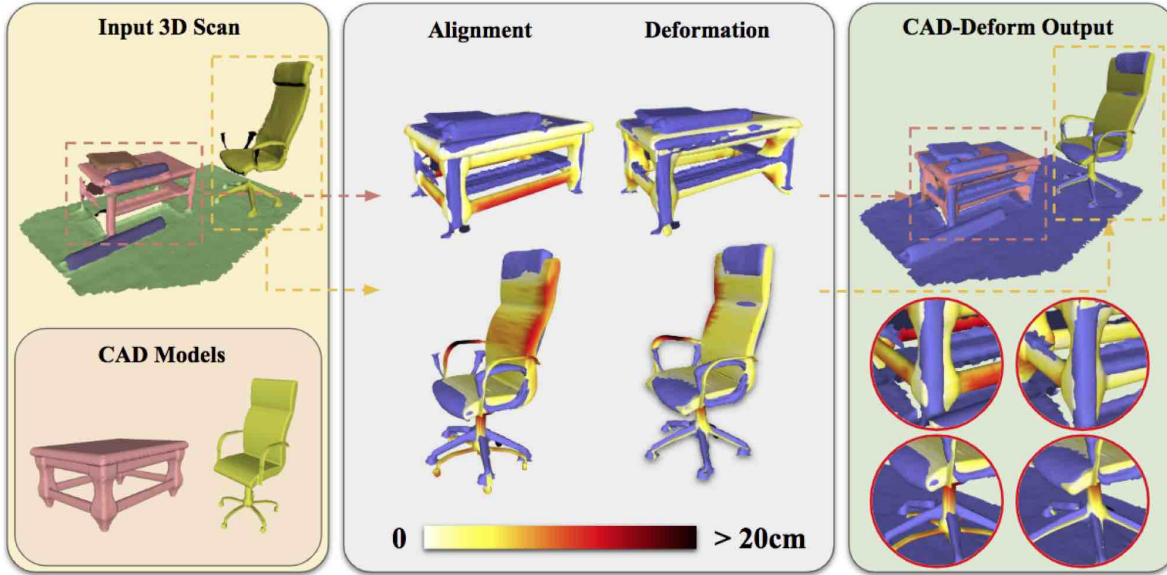


Figure 1: CAD-Deform takes as input a set of 3D CAD models aligned on a RGB-D scan (left). In order to achieve tight fits (middle), we propose a novel part-based deformation formulation that maintains the desired CAD properties such as sharp features.

An approach to overcome these problems has been proposed in [7, 8] and more recently developed using modern ML methods in [9, 10]. Building on the availability of parametric (CAD) models [11, 12] they perform reconstruction by finding and aligning similar CAD models from a database to each object in a noisy scan. To realize this approach, the authors of [9] introduced the Scan2CAD dataset comprising of pairwise keypoint correspondences and 9 DoF (degrees of freedom) alignments between instances of unique synthetic models from ShapeNet [11] and reconstructed scans from ScanNet [13]; in order to find and align CAD models to an input scan, they developed a deep neural model to predict correspondences, with a further optimization over potential matching correspondences for each candidate CAD model. The difference of [10] compared to the approach from [9] is an end-to-end procedure, combining initially decoupled steps to take into account additional feedback through the pipeline by learning the correspondences specifically tailored for the final alignment task.

However, geometric fidelity achieved between scans and CAD objects remains limited. CAD model geometry (clean and complete) differs significantly from scan geometry in low-level geometric features. As these methods focus on finding alignments by optimizing a small number of parameters (9 DoF), resulting alignments only roughly approximate scans, not capturing geometric details such as variation in 3D shapes of parts of individual objects.

In contrast, to improve geometric fidelity while keeping the benefit of mesh-based representations, we propose to increase the number of degrees of freedom by allowing the CAD objects to *deform* rather than stay rigid. In this work, we introduce a deformation framework CAD-Deform, which significantly increases the geometric quality of object alignment regardless of the alignment scheme. For an input scan, given object location and 9 DoF alignment of a potentially matching CAD model, we apply a specially designed mesh deformation procedure resulting in a more accurate shape representation of the underlying object. The deformation matches each semantic part of a 3D shape extracted from the PartNet [14] to the corresponding data in the scans and keeps sufficient rigidity and smoothness to produce perceptually acceptable results while minimizing the distance to scanned points. Thus, even if the initial location and alignment are not entirely accurate, the deformation can compensate to a significant extent for the discrepancy.

Our approach builds highly detailed scene descriptions with a high level of semantic accuracy for applications in 3D graphics. The approach outperforms state-of-the-art methods for CAD model alignment and mesh deformation by 2.1–6.3% for real-world 3D scans. To the best of our knowledge, the approach we propose is the first to use mesh deformation for scan-to-CAD alignment and real-world scene reconstruction. In summary, our work has several contributions:

- We developed a mesh deformation approach that 1) is computationally efficient, 2) does not require exact correspondences between each candidate CAD model and input scan, and 3) provides perceptually plausible

deformation thanks to a specially introduced smoothness term and inclusion of geometric features of a CAD model in the optimization functional.

- We developed a methodology to assess the fitting accuracy and the perceptual quality of the scan-to-CAD reconstruction. The methodology includes standard data fitting criteria similar to Chamfer distance to evaluate alignment accuracy, complimentary local and global criteria for visual quality assessment of resulting deformations, and a user study.
- We performed an ablation study to assess the influence of inaccuracies in the initial object location and alignment on the final reconstruction results. For that we used both ground-truth alignments from Scan2CAD dataset [9] along with predictions of their method, and alignments trained in the end-to-end fashion [10]. We compared results with the state-of-the-art methods for mesh deformation to highlight the advantages of our approach.

2 Related work

RGB-D scanning and reconstruction. RGB-D scanning and reconstruction are increasingly widely used, due to the availability of commodity range sensors and can be done both in real-time and offline modes. There are many methods for RGB-D-based real-time reconstruction such as KinectFusion [1], Voxel Hashing [3] or Bundle Fusion [6] that use the well-known volumetric fusion approach from [15]. ElasticFusion [4] is a representative offline approach to the reconstruction. These methods can produce remarkable results for large 3D environments. However, due to occlusions and imperfections of existing sensors, the reconstructed scenes contain many artifacts such as noise, missing surface parts, or over-smooth surfaces. Some methods aim to predict unobserved or corrupted parts of 3D scans from depth data. In [16], the world is modeled as a grid of voxels representing the signed distance to the nearest surface. Then structured Random Forest is used to predict the value of the signed distance function for each of the voxels computed to form a final occupancy estimate for unobserved voxels. [17] is to encode a depth map as a 3D volume and then aggregate both local geometric and contextual information via a 3D CNN to produce the probability distribution of voxel occupancy and object categories for all voxels inside the camera view frustum. Another approach is [18], where a 3D CNN architecture predicts a distance field from a partially-scanned input and finds the closest 3D CAD model from a shape database. By copying voxels from the nearest shape neighbors, they construct a high-resolution output from the low-resolution predictions, hierarchically synthesize a higher-resolution voxel output, and extracts the mesh from the implicit distance field. However, although all these methods can complete partial scans and improve 3D geometry reconstruction, the quality of the results still is far from artist-created 3D content.

CAD model alignment. Instead of reconstructing 3D geometry in a bottom-up manner, one can perform reconstruction by retrieving CAD models from a dataset and aligning them to the noisy scans. Matching CAD models to scans requires extracting 3D feature descriptors; thus, approaches have been proposed for 3D feature extraction. Hand-crafted features are often based on various histograms of local characteristics (*e.g.*, [19]). Such approaches do not generalize well to inherent variability and artifacts in real-world scans. Deep learning approaches lead to further improvements: *e.g.*, [20] propose a view consistency loss for a 3D keypoint prediction network based on RGB-D data; [21] develop 3D local learnable features for pairwise registration. After the descriptors are extracted, one can use a standard pipeline for CAD-to-scan alignment: first, matches between points based on 3D descriptors are identified, and then variational alignment is used to compute 6- or 9-DoF CAD model alignments. Typical examples, realizing this two-step approach, are described in [7, 8, 22, 9, 10]. The most recent ones [9, 10] use learnable descriptors and differ in that the latter considers an end-to-end scan-to-CAD alignment, reconstructing a scene in a single forward pass. An obvious limitation of these two-step pipelines is that the resulting alignments approximate scans only coarsely due to a pre-defined set of available CAD models and a highly constrained range of transformations. Other approaches in this category [23, 24, 25], although relying on the same two-step strategy, use only a single RGB or RGB-D input.

Mesh deformation. To improve surface reconstruction accuracy and obtain a more faithful description of geometric details (*e.g.*, preserve distinctive geometric features), it is desirable to consider a more general space of deformations than a simple non-uniform scaling. To this end, a mesh deformation approach based on Laplacian surface editing is presented in [26]. Another iterative mesh deformation scheme [27] imposes rigidity on the local transformations. Despite their conceptual simplicity, both methods require specifying correspondences between mesh vertices and scan data points. The same is true for [28, 29, 30] and unsuitable in our setting, due to extremely low (2–8) number of correspondences available per part, that additionally may not even be well defined for noisy and incomplete real-world scans. Many methods [29, 31, 32, 33] focus on automatic posing of human models, dense surface tracking and similar applications. However, while producing compelling results, the methods implicitly use the assumption that a template mesh and its target shape are very similar; as a result, the semantic parts of individual 3D objects are not changed either in shape or relative scale. In our work, we are comparing to ARAP (as-rigid-as-possible) [27] and Harmonic [34, 35] mesh deformation methods, however, with an added Laplacian smoothing term to leverage second-order information




Figure 2: Data acquisition for CAD-Deform: we project PartNet labels onto aligned ShapeNet CAD models (left), register these models to a 3D scene scan (middle), and extract points on the scene within ε -neighborhood of aligned mesh surface (we set $\varepsilon = 10$ cm), copying labels corresponding to nearest part of CAD model (right).

about the surface. This modification makes ARAP/Harmonic similar to [29, 28] as far as non-data-dependent energy terms are concerned. Other methods exist that propose non-linear constraints [36, 37, 38]. The energy terms of our framework were designed as a natural match for our problem: we define local 3D transformations on the CAD model, mapping each subpart to the 3D scene volume, and require smooth changes of these transformations over the model. In contrast to most deformation methods, we do not aim to keep the surface close to isometric: *e.g.*, a table reference model can be stretched highly anisotropically to match a different table in the data. Our energy is defined using local 3D affine transform primitives, penalizing sharp changes, without penalizing anisotropic deformations. These primitives also allow us to express 1D feature preservation simply, and the non-data terms are quadratic which is critical for our efficient preconditioned optimization. Methods in [36, 37, 38] propose non-linear energies, focusing on large rotations, but implicitly assuming quasi-isometry; adapting these methods is nontrivial.

3 Overview of CAD-Deform framework

Our approach is built on top of the framework from [9, 10] for CAD model retrieval and 9 DoF alignment in 3D scans. By running any of the approaches from [9, 10] for an input scan, we obtain initial object locations and 9 DoF alignments of CAD models potentially matching specific parts of the scan. Next, we apply our proposed mesh deformation procedure (see Section 4), resulting in a more accurate shape representation of the aligned objects:

1. We segment the CAD models into semantic 3D parts following the labelling from the PartNet dataset [14].
2. For each aligned object, we select points in the scan that are the nearest (within some fixed radius) to each vertex of the CAD model. We assign a label of the nearest part of the aligned CAD model to each such point.
3. As an input to the proposed mesh deformation procedure, we use the mesh model with semantic part labels and labelled segment of the 3D scene.
4. We deform the mesh by optimizing the energy depending on the relative positions of mesh vertices and labelled points of the scene, see Section 4.

4 Data-driven shape deformation

In this section, we describe our method for fitting a closest-match mesh from the shape dataset to the scanned data. Our algorithm assumes two inputs:

- an initial mesh \mathbf{M} from the dataset, with part labels assigned to vertices and 9 degrees-of-freedom (9 DoF) transformation assigned to mesh;
- a subset of the scanned data, segmented with the same part labels as \mathbf{M} .

Fig. 2 shows an example of the input data, see also Section 3.

Notations. A mesh $\mathbf{M} = (\mathbf{V}^0, \mathbf{E}, \mathbf{F})$ consists of a set of initial 3D vertex coordinates \mathbf{V}^0 of size n_v , the edge set \mathbf{E} of size n_e , and the triangle face set \mathbf{F} of size n_f . We compute the deformed vertex positions \mathbf{V} by minimizing a deformation energy. All vertices are assigned with part labels, which is a map $\mathbf{Q} : \mathbf{V}^0 \rightarrow \mathbf{C}$, where \mathbf{C} is the set of labels $c_i, i = 1, \dots, n_c$, and n_c is the number of parts in all objects in our dataset. For a mesh \mathbf{M}_m , $\mathbf{C}_m \subset \mathbf{C}$ is the set of labels of its parts. The set of n_p points \mathbf{P} has the same labels as the mesh \mathbf{M}_m we fit, *i.e.*, $\mathbf{C}_p = \mathbf{C}_m$. In addition, we assume that for every mesh \mathbf{M}_m on the scan we have a scaling transformation, represented by a 4×4 matrix T_m^0 , that aligns it with voxelized points \mathbf{P} . In our optimization, we use *per-edge* transformations T_e , discussed below, which we use to measure the deviation from a scaled version of the original shape and deformation smoothness.

4.1 Deformation energy

Our goal is to define a deformation energy to match the scanned data as closely as possible, while maintaining continuity and deformation smoothness, and penalizing deviation from the original shape. In addition, we include a term that preserves perceptually important linear geometric features of the mesh.

Conceptually, we follow common mesh deformation methods such as ARAP [27] which estimate a local transformation from vertex positions and penalize the deviation of this transformation from the target (*e.g.*, closest rotation). An important distinction is that the transformation we need to estimate locally is a 3D, rather than 2D transformation, and cannot be estimated from the deformed positions of vertices of a single triangle.

Edge transformations. We associate local 3D affine transformations with mesh edges. Each transformation is defined in a standard way by a 4×4 matrix in projective coordinates, with the last row $(0, 0, 0, c)$, $c \neq 0$. The vertices of two faces (i_1, i_2, i_3) and (i_2, i_1, i_4) , incident at an edge $e = (i_1, i_2)$, form a (possibly degenerate) tetrahedron. If it is not degenerate, then there is a unique linear transformation T_e mapping the undeformed positions $(v_{i_1}^0, v_{i_2}^0, v_{i_3}^0, v_{i_4}^0)$ to $(v_{i_1}, v_{i_2}, v_{i_3}, v_{i_4})$. Matrix T^0 of the affine transformation has 12 coefficients that are uniquely determined by the equations $T_e v_i^0 = v_i$, $i = 1 \dots 4$; moreover, these are linear functions of the deformed positions v_i , as these are only present on the right-hand side of the equations. Handling of degenerate tetrahedra is discussed in Section 4.2.

Energy. We define the following non-linear objective for the unknown deformed vertex positions \mathbf{V} and (a fixed) point set \mathbf{P} :

$$\begin{aligned} \mathcal{E}(\mathbf{V}, \mathbf{P}) &= \underbrace{E_{\text{shape}} + \alpha_{\text{smooth}} E_{\text{smooth}} + \alpha_{\text{sharp}} E_{\text{sharp}}}_{\text{quadratic problem}} + \alpha_{\text{data}} E_{\text{data}}, \\ E_{\text{shape}} &= \underbrace{\sum_{e \in \mathbf{E}} \|T_e(\mathbf{V}) - T_e^0\|_2^2}_{\text{deviation}}; & E_{\text{smooth}} &= \sum_{f \in \mathbf{F}} \sum_{e_i, e_j \in f} \|T_{e_i}(\mathbf{V}) - T_{e_j}(\mathbf{V})\|_2^2; \\ E_{\text{sharp}} &= \sum_{k=1}^{n_p} \sum_{e_s \in \mathbf{E}_{\text{sharp}}^k} \|T_{e_s}(\mathbf{V}) - T_{e_{s+1}}(\mathbf{V})\|_2^2; & E_{\text{data}} &= f_{\text{data}}(\mathbf{V}, \mathbf{P}). \end{aligned} \quad (1)$$

The first term penalizes deviations of the 3D affine transformations defined by the deformed vertex positions from the transformation (a non-uniform scale) that aligns mesh with the data. This term directly attempts to preserve the shape of the object, modulo rescaling.

As explained above, the transformations $T_e(\mathbf{V})$ are defined for non-degenerate input mesh configurations. Suppose the four initial vertex positions $(v_{i_1}^0, v_{i_2}^0, v_{i_3}^0, v_{i_4}^0)$ form a degenerate tetrahedron \mathbf{W} , *i.e.*, two faces incident at the edge are close to co-planar. In this case, we use an energy term consisting of two terms defined per triangle. Instead of 4×4 matrix for each non-degenerate tetrahedron, there is a 3×2 matrix for each triangular face of degenerate tetrahedron that represents a transformation restricted to the plane of this face. Note that in this case, the deformation in the direction perpendicular to the common plane of the triangles does not affect the energy, as it does not have an impact on the local shape of the deformed surface. We explain the remaining energy terms in the next sections.

4.2 Quadratic terms

Smoothness term. This term can be thought of as a discrete edge-based Laplacian, applied to the transforms associated with edges: the difference of any pair of transforms for edges belonging to the same triangle is penalized.

Sharp features term. We have observed that a simple way to preserve some of the perceptually critical geometric aspects of the input meshes is to penalize the change of deformations T_e along sharp geometric features, effectively preserving their overall shape (although still allowing, possibly non-uniform, scaling, rotations and translations). We detect sharp edges based on a simple dihedral angle thresholding, as this is adequate for the classes of CAD models we use. Detected sharp edges are concatenated in sequences, each consisting of vertices with exactly two incident sharp edges each, except the first and the last. Those sequences are defined for each part of the mesh. Sequences of different parts has no common sharp edges or vertices belonging to them. Effectively, the sharpness term increases the weights for some edges in the smoothness term.

4.3 Data term

We use two approaches to defining the data term, one based on screened attraction between close points in the mesh \mathbf{M}_m and \mathbf{P} , and the other one based on attraction between *a priori* chosen corresponding points of the mesh \mathbf{M}_m and data points in the set \mathbf{P} . We found that the former method works better globally, when the deformed mesh is still far from the target point cloud, while the latter is able to achieve a better match once a closer approximation is obtained.

Part-to-part mapping. We define a data-fitting term based on point proximity: we set an energy proportional to the distance between sufficiently close points. To avoid clustering of mesh points, we add a screening term that disables attraction for mesh vertices close to a given point.

We denote by $H(x)$ the Heaviside function, *i.e.*, a function with value 1 if $x \geq 0$ and zero otherwise. We set

$$f_{\text{data}}^{\text{p2p}}(\mathbf{V}, \mathbf{P}) = \sum_{c \in \mathbf{C}} \sum_{v \in \mathbf{V}_c} \sum_{p \in \mathbf{P}_c} \xi^\sigma(p, \mathbf{V}_c) (d^\varepsilon(v - p))^2, \quad (2)$$

where $d^\varepsilon(v - p) = (v - p) \cdot H(\varepsilon - \|v - p\|_2^2)$, and $\xi^\sigma(p, \mathbf{V}) = H(\min_{\{v \in \mathbf{V}\}} \|v - p\| - \sigma)$ is a ‘‘screening’’ function. The value for σ is chosen to be the mean edge length, the value of ε is about 10 edge lengths. To make $f_{\text{data}}^{\text{p2p}}$ differentiable, instead of the Heaviside function and min, we can use their smoothed approximations.

Nearest-neighbor mapping. Recall that the vertices $\{v_i\}$ of the mesh m and the points $\{p_i\}$ of the set \mathbf{P} have the same part label sets $\mathbf{C}_m = \mathbf{C}_\mathbf{P}$. For each label c , we consider \mathbf{V}_c^0 and \mathbf{P}_c , the set of mesh vertices in initial positions and the set of points with the same label c in \mathbf{P} . Let $B_{\mathbf{V}_c^0}, B_{\mathbf{P}_c}$ be the bounding boxes of these sets, and consider the affine transform T_B^c that maps $B_{\mathbf{V}_c^0}$ to $B_{\mathbf{P}_c}$. Among all possible correspondences between corners of the boxes, we choose the one that produces the affine transform closest to identity. Then the index $i = \iota^T(p)$ of the vertex v_i , corresponding to a point $p \in \mathbf{P}_c$, is determined as $\iota^c(p) = \arg \min_{\{i, v_i \in \mathbf{V}_c\}} \|T_B^c v_i - p\|_2^2$. Then the data term is defined as

$$f_{\text{data}}^{\text{nn}}(\mathbf{V}, \mathbf{P}) = \sum_{c \in \mathbf{C}} \sum_{p \in \mathbf{P}_c} \|p - v_{\iota^c(p)}\|_2^2. \quad (3)$$

4.4 Optimization

The optimization of (1) is highly nonlinear due to the data term. However, all other terms are quadratic functions of vertex coordinates, so minimizing

$$\begin{aligned} E_{\text{quad}} &= E_{\text{shape}} + \alpha_{\text{smooth}} E_{\text{smooth}} + \alpha_{\text{sharp}} E_{\text{sharp}} \\ &= \bar{\mathbf{V}}^\top A_{\text{shape}} \bar{\mathbf{V}} + \alpha_{\text{smooth}} \bar{\mathbf{V}}^\top A_{\text{smooth}} \bar{\mathbf{V}} + \alpha_{\text{sharp}} \bar{\mathbf{V}}^\top A_{\text{sharp}} \bar{\mathbf{V}} + b^\top \bar{\mathbf{V}} \end{aligned}$$

is equivalent to solving a system of linear equations. Denoting the sum of the matrices in the equation by A_{quad} , we obtain the optimum by solving $A_{\text{quad}} \bar{\mathbf{V}} = 0$ where the vector $\bar{\mathbf{V}}$ is a flattening of the vector \mathbf{V} of 3D vertex positions to a vector of length $3n_{\mathbf{V}}$.

The data term is highly nonlinear, but solving the complete optimization problem can be done efficiently using A_{quad}^{-1} as the preconditioner. For our problem, we use the preconditioned L-BFGS optimizer summarized in Algorithm 1.

Algorithm 1: Preconditioned L-BFGS mesh optimization (PL-BFGS)

$M_{\text{precond}} = A_{\text{quad}}^{-1}$ // stored as LU decomposition

$\bar{\mathbf{V}} = T_m^0(\bar{\mathbf{V}}^0)$

for $i \leftarrow 0$ **to** N_{iter} **do**

$g_{\text{tot}} = \alpha_{\text{data}} dE_{\text{data}}/dp + A_{\text{quad}} \bar{\mathbf{V}} + b$
 $\bar{\mathbf{V}} = \text{L-BFGS-step}(\bar{\mathbf{V}}, g_{\text{tot}}, M_{\text{precond}})$

end

5 Datasets

Our method relies on a number of publicly available datasets to compute mesh annotations for our deformations and assess fitting performance. To assess fitting performance, we use Scan2CAD dataset [9] that consists of 14225 ground-truth 9 DoF transformations between objects in 1506 reconstructions of indoor scenes from Scannet [13] and 3049 unique CAD models from ShapeNet [11].

Our deformation framework requires high-quality watertight meshes to support numerically stable optimization, which does not hold for ShapeNet CAD models. Thus, we remesh these to around 10k–15k vertices using [39], obtaining more uniform discretizations. To annotate these remeshed CAD models with semantic part labels required by our deformation procedure, we register them with the corresponding part-labeled meshes from the PartNet dataset [14] by first choosing an appropriate 90° rotation around each axis and then optimizing for a finer alignment using a point-to-point ICP algorithm [40] between vertices of the two meshes. We annotate the semantic parts for vertices in each mesh by projecting it from the respective closest vertices of the registered PartNet mesh.

The original ShapeNet meshes, however, are used to extract sharp geometric features, as these have easily detectable sharp angles between adjacent faces. We label as sharp all edges adjacent to faces with a dihedral angle smaller than the threshold $\alpha_{\text{sharp}} = 120^\circ$. We further project vertex-wise sharpness labels from the original to the remeshed CAD models and select a sequence of edges forming the shortest paths between each pair of vertices as sharp.

6 Results

6.1 Evaluation setup

Our performance evaluation of obtained deformations is multifaceted and addresses the following quality-related questions:

- Scan fitting performance: *How well do CAD deformations fit?*
- Perceptual performance for deformations: *How CAD-like are deformations?*
- Contributions of individual energy terms: *Which energy terms are essential?*
- Deformation flexibility: *Can better shapes be achieved by approximating clean meshes rather than noisy scans?*

Fitting and perceptual metrics. We quantify the deformation performance in terms of fitting quality between the scene scans and 3D CAD models using a family of related measures computed on a per-instance basis. For vertices $\mathbf{V} = (v_i)$ of the deformed mesh \mathbf{M} , we compute distances to their respective nearest neighbors ($\text{NN}(v_i, \mathbf{S})$) in the scan \mathbf{S} . We compute per-instance Accuracy = $|\mathbf{V}_{\text{close}}|/|\mathbf{V}|$, reflecting the fraction of closely located vertices, and trimmed minimum matching distance $\text{tMMD} = \sum_{v_i \in \mathbf{V}} \min(\tau, \|v_i - \text{NN}(v_i, \mathbf{S})\|_1)/|\mathbf{V}|$, where $\mathbf{V}_{\text{close}} = \{v_i \in \mathbf{V} : \|v_i - \text{NN}(v_i, \mathbf{S})\|_1 < \tau\}$

is the set of vertices falling within L_1 -distance τ to their nearest neighbor in the scan, and τ controls robustness w.r.t. incomplete scans. We set $\tau = 0.2$ (see supplementary) in our experiments and report Accuracy and tMMD values averaged over classes and instances.

There is no universally agreed perceptual quality measure for meshes; thus, we opted for a tripartite evaluation for our resulting deformations. First, we measure dihedral angle mesh error (DAME) [41], revealing differences in *local*

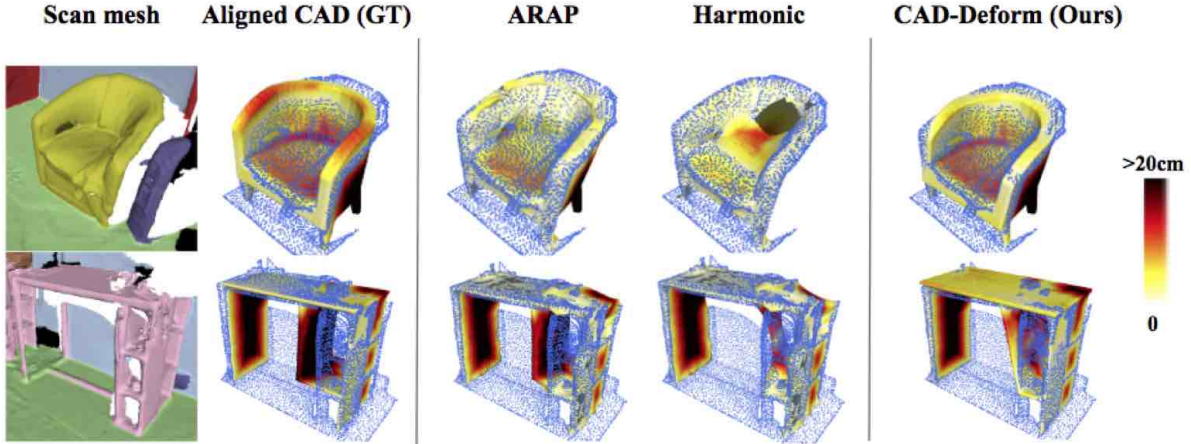


Figure 3: Deformations obtained using our method and the baselines, with mesh colored according to the Euclidean distance to its nearest point in the scan. We note that high accuracy scores for Harmonic and ARAP deformations are achieved at the cost of destroying the initial structure of the mesh, particularly in regions where scan is missing (note that back side and armrests are gone for chairs in the first and second rows). In contrast, our method is better able to preserve smooth surfaces, sharp features, and overall mesh integrity, while keeping accurate local alignment.

surface quality between the original and the distorted meshes:

$$\text{DAME}(\mathbf{M}, \mathbf{M}_{\text{def}}) = \frac{1}{|\mathbf{E}|} \sum_{\text{adjacent } f_1, f_2} |D_{f_1, f_2} - \overline{D_{f_1, f_2}}| \cdot \exp \left\{ (Z_{\text{DAME}} D_{f_1, f_2})^2 \right\},$$

where D_{f_1, f_2} and $\overline{D_{f_1, f_2}}$ represent oriented dihedral angles between faces f_1 and f_2 in the original and deformed meshes, respectively, and $Z_{\text{DAME}} = \sqrt{\log(100/\pi)}/\pi$ is a parameter scaling DAME values to $[0, 100]$.

Second, we assess *abnormality of deformed shapes* with respect to the distribution of the undeformed shapes, building on the idea of employing deep autoencoders for anomaly detection in structured high-dimensional data. By replicating the training instances, autoencoders learn features that minimize reconstruction error; for novel instances *similar* to those in the training set, reconstruction error is low compared to that of strong *outliers*. We train six autoencoders [42, 43] for point clouds using vertices of undeformed meshes separately for the top six classes present in Scan2CAD annotation: *table*, *chair*, *display*, *trashbin*, *cabinet*, and *bookshelf*. Passing vertices \mathbf{V}_{def} of a deformed shape to the respective autoencoder, one can assess how accurately deformed meshes can be approximated using features of undeformed meshes. This property can be evaluated with Earth Mover’s Distance (EMD) $d_{\text{EMD}}(\mathbf{V}_{\text{def}}, \mathbf{V}'_{\text{def}}) = \min_{\phi: \mathbf{V}_{\text{def}} \rightarrow \mathbf{V}'_{\text{def}}} \sum_{v \in \mathbf{V}_{\text{def}}} \|v - \phi(v)\|_2$,

where ϕ is a bijection, obtained as a solution to the optimal transportation problem involving \mathbf{V}_{def} and \mathbf{V}'_{def} , that can intuitively be viewed as the least amount of work needed to transport \mathbf{V}_{def} vertices to positions of \mathbf{V}'_{def} .

Lastly, we assess *real human perception* of deformations in a user study, detailed in Section 6.3.

Optimization details. To perform quantitative comparisons, we use 299 scenes in ScanNet constituting Scan2CAD validation set [9], but with 697 shapes present in PartNet dataset, amounting to 1410 object instances. Our full experimental pipeline is a sequence of deformation stages with different optimization parameters, and Hessian being recomputed before each stage. Specifically, we perform one *part-to-part* optimization with parameters $\alpha_{\text{shape}} = 1$, $\alpha_{\text{smooth}} = 0$, $\alpha_{\text{sharp}} = 0$, $\alpha_{\text{data}} = 5 \times 10^4$ for 100 iterations, then we perform 5 runs of *nearest-neighbor* deformation for 50 iterations with parameters $\alpha_{\text{shape}} = 1$, $\alpha_{\text{smooth}} = 10$, $\alpha_{\text{sharp}} = 10$, $\alpha_{\text{data}} = 10^3$. More details about optimization and timings are provided in the supplementary.

6.2 Fitting Accuracy: How well do CAD Deformations fit?

We first demonstrate how deformation affects scan fitting performance for meshes aligned using different methods, specifically, we use true-positive shape alignments computed using Scan2CAD (S2C) [9], End-to-End (E2E) [9], as well as ground-truth alignments. We start with an aligned mesh, copy the 9 DoF transformation to each of the mesh vertices, and optimize using our deformation method with parameters described in Section 6.1. We report Accuracy

| Method | Class avg. | | | Instance avg. | | |
|------------------|-------------|-------------|-------------|---------------|-------------|-------------|
| | GT | S2C [9] | E2E [10] | GT | S2C [9] | E2E [10] |
| # TPs | 1410 | 499 | 882 | 1410 | 499 | 882 |
| TP undeformed | 89.2 | 83.7 | 88.5 | 90.6 | 79.4 | 93.9 |
| Ours: NN only | 89.7 | 84.3 | 89.0 | 91.4 | 84.7 | 94.4 |
| Ours: p2p only | 90.3 | 88.3 | 89.4 | 91.6 | 90.3 | 94.9 |
| Ours: w/o smooth | 90.6 | 90.0 | 89.6 | 92.3 | 90.3 | 95.0 |
| Ours: w/o sharp | 90.3 | 86.9 | 90.6 | 92.3 | 89.4 | 95.2 |
| CAD-Deform | 91.7 | 89.8 | 90.3 | 93.1 | 92.8 | 94.6 |

Table 1: Comparative evaluation of our deformations to true positive (TP) alignments by non-deformable approaches in terms of Accuracy (%). Note that deformations improve performance for all considered alignment approaches.

| Method | bookshelf | cabinet | chair | display | table | trashbin | other | class avg. | avg. |
|--------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| # instances | 142 | 162 | 322 | 86 | 332 | 169 | 197 | 201.4 | 1410 |
| Ground-truth | 88.0 | 75.2 | 94.8 | 98.9 | 89.6 | 96.6 | 81.4 | 89.2 | 90.6 |
| Ours | 90.5 | 82.2 | 95.4 | 99.1 | 91.0 | 98.6 | 84.8 | 91.7 | 93.1 |

Table 2: Comparative evaluation of our approach to non-deformable ground-truth and baselines in terms of scan approximation Accuracy (%). We conclude that our deformations improve fitting accuracy across all object classes by 2.5 % on average.

scores in terms of fraction of well-approximated points in the scan for aligned shapes pre- and post-optimization in Table 2, achieving improved performance across all considered alignment procedures.

Surprisingly, we improve even over ground-truth alignments by as much as 2.5 %. Thus, we compute per-class scores in Table 1 (comparison across alignments), reporting improvements of up to 7 % in average scan approximation accuracy that are consistent *across all object classes*. We visualize example deformations obtained using our approach and baselines in Figure 3.

We have discovered our deformation framework to be robust w.r.t. level of detail in the data term in (2) and provide more detail in the supplementary.

6.3 CAD Quality: How CAD-like are deformed models?

Having obtained a collection of deformed meshes, we aim to assess their visual quality in comparison to two baseline deformation methods: as-rigid-as-possible (ARAP) [27] and Harmonic deformation [34, 35], using a set of perceptual quality measures. To bring second-order information about mesh surface in energy formulation of ARAP/Harmonic, we add the Laplacian smoothness term $E_{\text{Lap}}(\mathbf{V}, \mathbf{V}') = \sum_{i=1}^{|\mathbf{V}|} \|L(v_i) - L(v'_i)\|^2$, where $L(v_i) = \frac{1}{|\mathcal{N}(i)|} \sum_{j=1}^{|\mathcal{N}(i)|} v_j$ and $\mathcal{N}(i)$ is a set of one-ring neighbors of vertex $v_i \in \mathbf{V}$ ($v'_i \in \mathbf{V}'$). The details of our user study design and visual assessment are provided in the supplementary.

| Method | DAME | | EMD $\times 10^{-3}$ | | User study | Accuracy | |
|-------------------|-------------|-------------|----------------------|-----------|------------|-------------|-------------|
| | cls. | inst. | cls. | inst. | | cls. | inst. |
| No deformation | 0 | 0 | 77 | 77 | 8.6 | 89.2 | 90.6 |
| ARAP [27] | 47.1 | 45.7 | 88 | 87 | 4.0 | 90.8 | 91.8 |
| Harmonic [34, 35] | 65.1 | 65.2 | 104 | 102 | 2.6 | 96.2 | 96.6 |
| CAD-Deform | 20.5 | 17.2 | 84 | 84 | 7.7 | 91.7 | 93.1 |

Table 3: Quantitative evaluation of visual quality of deformations obtained using ARAP [27], Harmonic deformation [34, 35], and our CAD-Deform, using a variety of local surface-based (DAME [41]), neural (EMD [42, 43]), and human measures.

6.4 Ablation study

To evaluate the impact of individual energy terms in (1) on both scan fitting performance and perceptual quality, we exclude each term from the energy and compute deformations by optimizing the remaining ones. First, we exclude sharpness or smoothness terms, optimizing for deformations using the original two-stage method; second, to better understand the influence of each stage, we perform experiments with only the first or the second stage (a single run only). We aggregate results into Table 1 and display them visually in Figure 4, concluding that our CAD-Deform maintains the right balance between fit to the scan and perceptual quality of resulting deformations.




Figure 4: Qualitative results of ablation study using our deformation framework, with mesh coloured according to the value of the tMMD measure. Note that including smoothness term is crucial to prevent surface self-intersections, while keeping sharpness allows to ensure consistency in parallel planes and edges.




Figure 5: Qualitative shape translation results, interpolating between the original mesh (left) and the target mesh (right).

6.5 Shape morphing results

To demonstrate the ability of our mesh deformation framework to perform shape interpolation, we choose two different meshes in the same ShapeNet category and optimize our energy (1) to approximate one with the other, see Fig. 5.

7 Conclusion

In this work, we have presented CAD-Deform, a shape deformation-based scene reconstruction method leveraging CAD collections that is capable of improving over existing alignment methods. More specifically, we introduce a composite deformation energy formulation that achieves regularization from semantic part structures, enforces smooth transformations, and preserves sharp geometric features. As a result we obtain significantly improved perceptual quality of final 3D CAD models compared to state-of-the-art deformation formulations, such as ARAP and polyharmonic deformation frameworks. Overall, we believe that our method is an important step towards obtaining lightweight digital replica from the real world that are both of high-quality and accurate fits at the same time.

Acknowledgements

The authors acknowledge the usage of the Skoltech CDISE HPC cluster Zhores for obtaining the results presented in this paper. The work was partially supported by the Russian Science Foundation under Grant 19-41-04109.

References

- [1] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. Kinectfusion: Real-time 3d reconstruction and interaction using a moving depth camera. In *UIST '11 Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 559–568. ACM, 2011.
- [2] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *IEEE ISMAR*. IEEE, October 2011.
- [3] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (TOG)*, 2013.
- [4] T. Whelan, S. Leutenegger, R. F. Salas-Moreno, B. Glocker, and A. J. Davison. ElasticFusion: Dense SLAM without a pose graph. In *Robotics: Science and Systems (RSS)*, Rome, Italy, July 2015.
- [5] Sungjoon Choi, Q. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5556–5565, June 2015.
- [6] Angela Dai, Matthias Nießner, Michael Zollöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface re-integration. *ACM Transactions on Graphics 2017 (TOG)*, 2017.
- [7] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H. J. Kelly, and Andrew J. Davison. Slam++: Simultaneous localisation and mapping at the level of objects. In *CVPR*, page 1352–1359. IEEE Computer Society, 2013.
- [8] Oliver Mattausch, Daniele Panozzo, Claudio Mura, Olga Sorkine-Hornung, and Renato Pajarola. Object detection and classification from large-scale cluttered indoor scans. In *Computer Graphics Forum*, volume 33, pages 11–21. Wiley Online Library, 2014.
- [9] Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2614–2623, 2019.
- [10] Armen Avetisyan, Angela Dai, and Matthias Nießner. End-to-end cad model retrieval and 9dof alignment in 3d scans. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2551–2560, 2019.
- [11] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [12] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [13] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017.
- [14] 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. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 909–918, 2019.
- [15] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '96*, page 303–312, New York, NY, USA, 1996. Association for Computing Machinery.
- [16] Michael Firman, Oisín Mac Aodha, Simon Julier, and Gabriel J Brostow. Structured prediction of unobserved voxels from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5431–5440, 2016.
- [17] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017.
- [18] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017.

- [19] Bertram Drost and Slobodan Ilic. 3d object detection and localization using multimodal point pair features. In *3DIMPVT*, pages 9–16. IEEE Computer Society, 2012.
- [20] Xingyi Zhou, Arjun Karapur, Chuang Gan, Linjie Luo, and Qixing Huang. Unsupervised domain adaptation for 3d keypoint estimation via view consistency. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 141–157, Cham, 2018. Springer International Publishing.
- [21] Haowen Deng, Tolga Birdal, and Slobodan Ilic. 3d local features for direct pairwise registration. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [22] Yangyan Li, Angela Dai, Leonidas Guibas, and Matthias Nießner. Database-assisted object retrieval for real-time 3d reconstruction. *Computer Graphics Forum*, 34(2):435–446, 2015.
- [23] Mathieu Aubry, Daniel Maturana, Alexei Efros, Bryan Russell, and Josef Sivic. Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *CVPR*, 2014.
- [24] Ruiqi Guo, Chuhan Zou, and Derek Hoiem. Predicting complete 3d models of indoor scenes. *arXiv preprint arXiv:1504.02437*, 2015.
- [25] Saurabh Gupta, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Aligning 3d models to rgb-d images of cluttered scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4731–4740, 2015.
- [26] Carsten Stoll, Zachi Karni, Christian Rössl, Hitoshi Yamauchi, and Hans-Peter Seidel. Template deformation for point cloud fitting. In *SPBG*, pages 27–35, 2006.
- [27] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116, 2007.
- [28] Miao Liao, Qing Zhang, Huamin Wang, Ruigang Yang, and Minglun Gong. Modeling deformable objects from a single depth camera. pages 167 – 174, 11 2009.
- [29] Jascha Achenbach, Eduard Zell, and Mario Botsch. Accurate face reconstruction through anisotropic fitting and eye correction. In *VMV*, pages 1–8, 2015.
- [30] B. Amberg, S. Romdhani, and T. Vetter. Optimal step nonrigid icp algorithms for surface registration. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.
- [31] Cedric Cagniart, Edmond Boyer, and Slobodan Ilic. Iterative mesh deformation for dense surface tracking. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1465–1472. IEEE, 2009.
- [32] Sang Il Park and Seong-Jae Lim. Template-based reconstruction of surface mesh animation from point cloud animation. *ETRI journal*, 36(6):1008–1015, 2014.
- [33] Tamal K Dey, Bo Fu, Huamin Wang, and Lei Wang. Automatic posing of a meshed human model using point clouds. *Computers & Graphics*, 46:14–24, 2015.
- [34] Mario Botsch and Leif Kobbelt. An intuitive framework for real-time freeform modeling. *ACM Transactions on Graphics (TOG)*, 23(3):630–634, 2004.
- [35] Alec Jacobson, Elif Tosun, Olga Sorkine, and Denis Zorin. Mixed finite elements for variational surface modeling. In *Computer graphics forum*, volume 29, pages 1565–1574. Wiley Online Library, 2010.
- [36] L. He and S. Schaefer. Mesh denoising via l0 minimization. *Proc. ACM SIGGRAPH*, pages 64:1–64:8, 01 2013.
- [37] Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '03*, page 62–67, Goslar, DEU, 2003. Eurographics Association.
- [38] Stefan Fröhlich and Mario Botsch. Example-driven deformations based on discrete shells. *Comput. Graph. Forum*, 30:2246–2257, 12 2011.
- [39] Jingwei Huang, Hao Su, and Leonidas Guibas. Robust watertight manifold surface generation method for shapenet models. *arXiv preprint arXiv:1802.01698*, 2018.
- [40] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. In *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, pages 145–152. IEEE, 2001.
- [41] Libor Váša and Jan Rus. Dihedral angle mesh error: a fast perception correlated distortion measure for fixed connectivity triangle meshes. *Computer Graphics Forum*, 31(5):1715–1724, 2012.
- [42] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International Conference on Machine Learning*, pages 40–49, 2018.

- [43] Vage Egiazarian, Savva Ignatyev, Alexey Artemov, Oleg Voynov, Andrey Kravchenko, Youyi Zheng, Luiz Velho, and Evgeny Burnaev. Latent-space laplacian pyramids for adversarial representation learning with 3d point clouds. 12 2019.

A Statistics on the used datasets

In Tables 4 & 5, we summarize statistical information on the number of instances and categories considered in our evaluation. As we require parts annotations as an important ingredient in our deformation, we only select instances in Scan2CAD [9] where the associated parts annotation in PartNet [14] is available, resulting in total in 9 categories (25%), 572 instances (18%), and 1979 annotated correspondences (14%). Note that the vast majority of cases remain within our consideration, keeping our evaluation comprehensive.

| Collection | Categories | Instances | Corresp. |
|---------------------|------------|-----------|----------|
| Scan2CAD [9] | 35 | 3,049 | 14,225 |
| w/parts annotations | 24 | 2,477 | 12,246 |

Table 4: Overall statistics on the numbers of categories, instances, and correspondences present in our datasets.

| Name | Scan2CAD | | PartNet \cap Scan2CAD | |
|---|----------|--------|-------------------------|--------|
| | corresp. | shapes | corresp. | shapes |
| <i>Shape categories used in our evaluation:</i> | | | | |
| chair | 4677 | 652 | 4351 | 632 |
| table | 2616 | 830 | 2594 | 822 |
| cabinet | 1401 | 310 | 1258 | 294 |
| trash bin | 1044 | 89 | 1042 | 88 |
| bookshelf | 824 | 150 | 812 | 145 |
| display | 770 | 165 | 762 | 161 |
| <i>Shape categories NOT used in our evaluation:</i> | | | | |
| bed | 355 | 50 | 342 | 47 |
| file cabinet | 294 | 70 | 290 | 68 |
| bag | 165 | 9 | 165 | 9 |
| lamp | 135 | 55 | 135 | 55 |
| bathtub | 474 | 96 | 129 | 25 |
| microwave | 99 | 37 | 98 | 36 |
| sofa | 577 | 247 | 60 | 20 |
| laptop | 51 | 24 | 51 | 24 |
| keyboard | 62 | 11 | 48 | 9 |

Table 5: The top 15 most frequent ShapeNet categories in Scan2CAD dataset including a detailed information on those with the availability of the corresponding parts annotations.

We further select the most well-presented six shape categories as our core evaluation set, outlined in Table 5. Note that as our method is non-learnable, we can just as easily experiment with the remaining categories, at the cost of somewhat reduced statistical power.

B Optimization details

Our full experimental pipeline is a sequence of deformation stages with different optimization parameters, and Hessian being recomputed before each stage. Specifically, we perform one *part-to-part* optimization with parameters $\alpha_{\text{shape}} = 1$, $\alpha_{\text{smooth}} = 0$, $\alpha_{\text{sharp}} = 0$, $\alpha_{\text{data}} = 5 \times 10^4$ for 100 iterations, then we perform 5 runs of *nearest-neighbor*

deformation for 50 iterations with parameters $\alpha_{\text{shape}} = 1, \alpha_{\text{smooth}} = 10, \alpha_{\text{sharp}} = 10, \alpha_{\text{data}} = 10^3$. Such number of iterations was sufficient to achieve convergence with energy changes less than 10^{-1} in our experiments. Runtime of our method breaks into cost computation (~ 0.3 s), backward (~ 0.2 s), and optimization steps containing the main bottleneck (sparse matrix-vector multiplication) (~ 1.2 s) for a typical 10^4 vertices mesh. All operations can be easily further optimized.

C Qualitative fitting results

| Method | bookshelf | cabinet | chair | display | table | trash bin | other | class avg. | avg. |
|------------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|
| ARAP | 52.48 | 41.77 | 45.52 | 51.30 | 41.77 | 57.00 | 39.75 | 47.08 | 45.67 |
| Harmonic | 64.77 | 58.74 | 68.06 | 64.22 | 58.26 | 80.13 | 61.70 | 65.12 | 65.18 |
| Ours: NN only | 21.54 | 23.39 | 7.31 | 18.37 | 18.69 | 18.13 | 16.07 | 17.64 | 14.14 |
| Ours: p2p only | 22.44 | 24.28 | 9.51 | 21.12 | 18.76 | 15.30 | 18.34 | 18.54 | 15.57 |
| Ours: w/o smooth | 27.15 | 29.27 | 14.50 | 27.05 | 24.48 | 24.39 | 23.26 | 24.30 | 20.95 |
| Ours: w/o sharp | 26.43 | 25.98 | 13.34 | 24.87 | 22.47 | 21.04 | 21.18 | 22.19 | 19.10 |
| CAD-Deform | 24.8 | 24.1 | 11.4 | 24.4 | 21.6 | 19.4 | 17.6 | 20.5 | 17.2 |

Table 6: Quantitative results of local surface quality evaluation using DAME measure [41] (the smaller, the better, normalized to a maximum score of 100), where our CAD-Deform compares favourably to the baselines across all considered categories. Note, however, how surface quality significantly decreases when smoothness and sharp feature-related terms are dropped.

| $\text{EMD} \times 10^{-3}$ | bookshelf | cabinet | chair | display | table | trash bin | class avg. | avg. |
|-----------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Ground-truth | 77.8 | 78.9 | 76.1 | 77.5 | 77.3 | 73.1 | 76.8 | 77.0 |
| ARAP [27] | 80.3 | 86.5 | 88.5 | 85.4 | 86.8 | 98.1 | 87.6 | 87.3 |
| Harmonic [34, 35] | 94.0 | 110.6 | 95.8 | 95.3 | 103.2 | 122.6 | 103.6 | 101.7 |
| CAD-Deform | 79.0 | 81.1 | 80.3 | 91.7 | 87.0 | 87.4 | 84.4 | 83.8 |

Table 7: Results of LSLP-GAN reconstruction in terms of Earth-Mover’s Distance between reconstructed and original point clouds of mesh vertices.

In Figure 6, we display a series of qualitative results with a variety of shape deformations with different classes of instances. Comparing to baselines, our framework achieves accurate fit while preserving sufficient perceptual quality.

Table 6 reports the results of surface quality evaluation using deformations obtained with our CAD-Deform vs. the baselines, category-wise. While outperforming the baseline methods across all categories, we discover the smoothness and sharpness energy terms to be the crucial ingredients in keeping high-quality meshes.

Figure 7 displays visually the deformation results using the three distinct classes, highlighting differences in surfaces obtained using the three methods.

Table 7 reports shape abnormality evaluation results across the six considered categories. Baselines show (Fig. 8) low reconstruction quality as evidenced by a larger number of black points. In other words, comparing to CAD-Deform, the distance from these meshes to undeformed ones is much larger.

In Figure 9, we show a series of examples for CAD-Deform ablation study. Perceptual quality degrades when excluding every term from the energy.

D Morphing

In this section, we present an additional series of examples of morphing properties (Fig. 10). Every iteration of optimization process gradually increases the quality of fit. With CAD-Deform we can morph each part to imitate the structure of the target shape.

| Accuracy, % | class avg. | avg. |
|--------------------|-------------------|--------------|
| Ground-truth | 89.22 | 90.56 |
| Level 1 (object) | 89.25 | 90.79 |
| Level 2 | 89.16 | 91.21 |
| Level 3 | 89.40 | 91.05 |
| Level 4 (parts) | 91.65 | 93.12 |

Table 8: Comparative evaluation of our approach in terms of Accuracy on different levels of detail.

E PartNet annotation

This set of experiments shows how quality of fitting depends on mesh vertices labelling. We can provide labels for mesh in different ways depending on the level in PartNet hierarchy [14]. We observe the increase of fitting quality with greater level of detail (Table 8). Examples presented in Figure 11 are selected as the most distinguishable deformations on different levels. There are minor visual differences in deformation performance of part labeling level.

F Fitting Accuracy analysis

CAD-Deform deformation framework is sensitive to Accuracy threshold τ for the distance between mesh vertices and close scan points. In Figure 12 variation of τ threshold is presented and we selected $\tau = 0.2$ m for fitting Accuracy metric.

G Perceptual assessment and user study details

Having obtained a collection of deformed meshes, we aim to assess their visual quality in comparison to two baseline deformation methods: as-rigid-as-possible (ARAP) [27] and Harmonic deformation [34, 35], using a set of perceptual quality measures. The details of our user study design and visual assessment are provided in the supplementary. To this end, we use original and deformed meshes to compute DAME and reconstruction errors, as outlined in Section 6.1, and complement these with visual quality scores obtained with a user study (see below). These scores, presented in Table 3, demonstrate that shapes obtained using CAD-Deform have $2\times$ higher surface quality, only slightly deviate from undeformed shapes as viewed by neural autoencoders, and receive $2\times$ higher ratings in human assessment, while sacrificing only 1.1–4.5 % accuracy compared to other deformation methods.

Design of our user study. The users were requested to examine renders of shapes from four different categories: the original undeformed shapes as well as shapes deformed using ARAP, Harmonic, and CAD-Deform methods, and give a score to each shape according to the following perceptual aspects: surface quality and smoothness, mesh symmetry, visual similarity to real-world objects, and overall consistency. Ten random shapes from each of the four categories have been rendered from eight different views and scored by 100 unique users on a scale from 1 (bad) to 10 (good). The resulting visual quality scores are computed by averaging over users and shapes in each category.

In Figure 13, we present a distribution of user scores over different deformation methods and shapes. It can be clearly seen that users prefer our deformation results to baselines for all of the cases, which is obvious from the gap between histogram of CAD-Deform and ARAP/Harmonic histograms. At the same time, shapes deformed by CAD-Deform are close to undeformed ShapeNet shapes in terms of surface quality and smoothness, mesh symmetry, visual similarity to real-world objects, and overall consistency. Besides, in Tables 9, 10 we provide numbers for evaluation of ARAP/Harmonic deformations w.r.t. the change of Laplacian term weight.

| Lap. term weight | Class avg. | | | Instance avg. | | |
|---------------------------------|------------|---------|----------|---------------|---------|----------|
| | GT | S2C [9] | E2E [10] | GT | S2C [9] | E2E [10] |
| $\alpha_{\text{Lap}} = 10^{-2}$ | 90.9 | 81.3 | 90.0 | 92.0 | 80.9 | 90.8 |
| $\alpha_{\text{Lap}} = 10^{-1}$ | 91.0 | 81.3 | 90.0 | 92.0 | 80.9 | 90.7 |
| $\alpha_{\text{Lap}} = 1$ | 91.0 | 81.3 | 89.9 | 91.9 | 80.9 | 90.7 |
| $\alpha_{\text{Lap}} = 5$ | 90.9 | 81.2 | 89.9 | 91.9 | 80.9 | 90.7 |
| $\alpha_{\text{Lap}} = 20$ | 90.9 | 81.2 | 89.9 | 91.8 | 80.8 | 90.6 |

Table 9: Comparative evaluation of ARAP deformations w.r.t. the change of Laplacian term weight in terms of Accuracy (%).

| Lap. term weight | Class avg. | | | Instance avg. | | |
|---------------------------------|------------|---------|----------|---------------|---------|----------|
| | GT | S2C [9] | E2E [10] | GT | S2C [9] | E2E [10] |
| $\alpha_{\text{Lap}} = 10^{-2}$ | 96.3 | 94.3 | 96.6 | 96.7 | 94.5 | 96.9 |
| $\alpha_{\text{Lap}} = 10^{-1}$ | 96.3 | 94.2 | 96.6 | 96.7 | 94.4 | 96.9 |
| $\alpha_{\text{Lap}} = 1$ | 96.3 | 94.2 | 96.6 | 96.7 | 94.2 | 96.9 |
| $\alpha_{\text{Lap}} = 5$ | 96.2 | 94.0 | 96.6 | 96.6 | 94.0 | 96.8 |
| $\alpha_{\text{Lap}} = 20$ | 96.2 | 93.8 | 96.5 | 96.6 | 93.8 | 96.7 |

Table 10: Comparative evaluation of Harmonic deformations w.r.t. the change of Laplacian term weight in terms of Accuracy (%).




Figure 6: Qualitative shape deformation results using obtained using ARAP [27], Harmonic deformation [34, 35], and our CAD-Deform. Mesh surface is colored according to the value of tMMD measure, with darker values corresponding to the larger distance values.

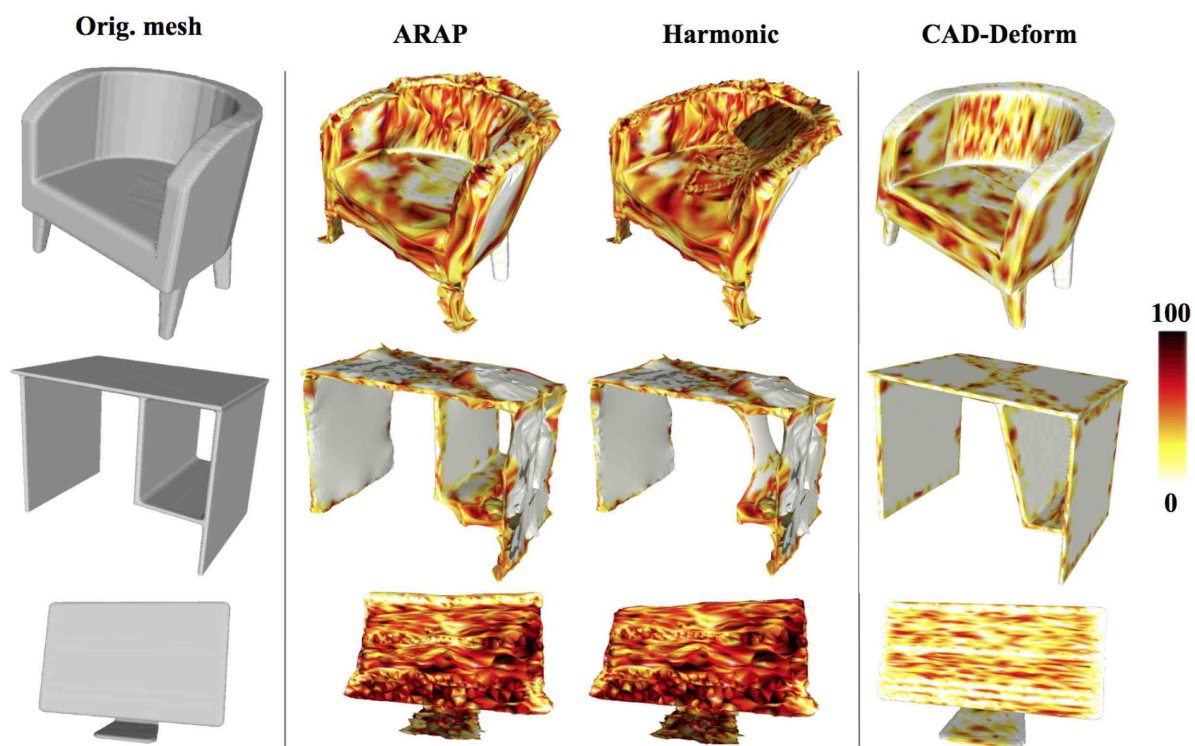


Figure 7: Qualitative comparison of deformations obtained using ARAP [27], Harmonic deformation [34, 35], and our CAD-Deform, with shapes coloured according to the value of DAME measure [41]. Our approach results in drastic improvements in local surface quality, producing higher-quality surfaces compared to other deformations.




Figure 8: Qualitative comparison of reconstruction of point clouds extracted from mesh vertices. These meshes are obtained using ARAP [27], Harmonic deformation [34, 35], and our CAD-Deform, the first column corresponds to original undeformed meshes. The color of reconstructed point clouds is related to Earth-Mover’s Distance between reconstructed and original point clouds of mesh vertices.

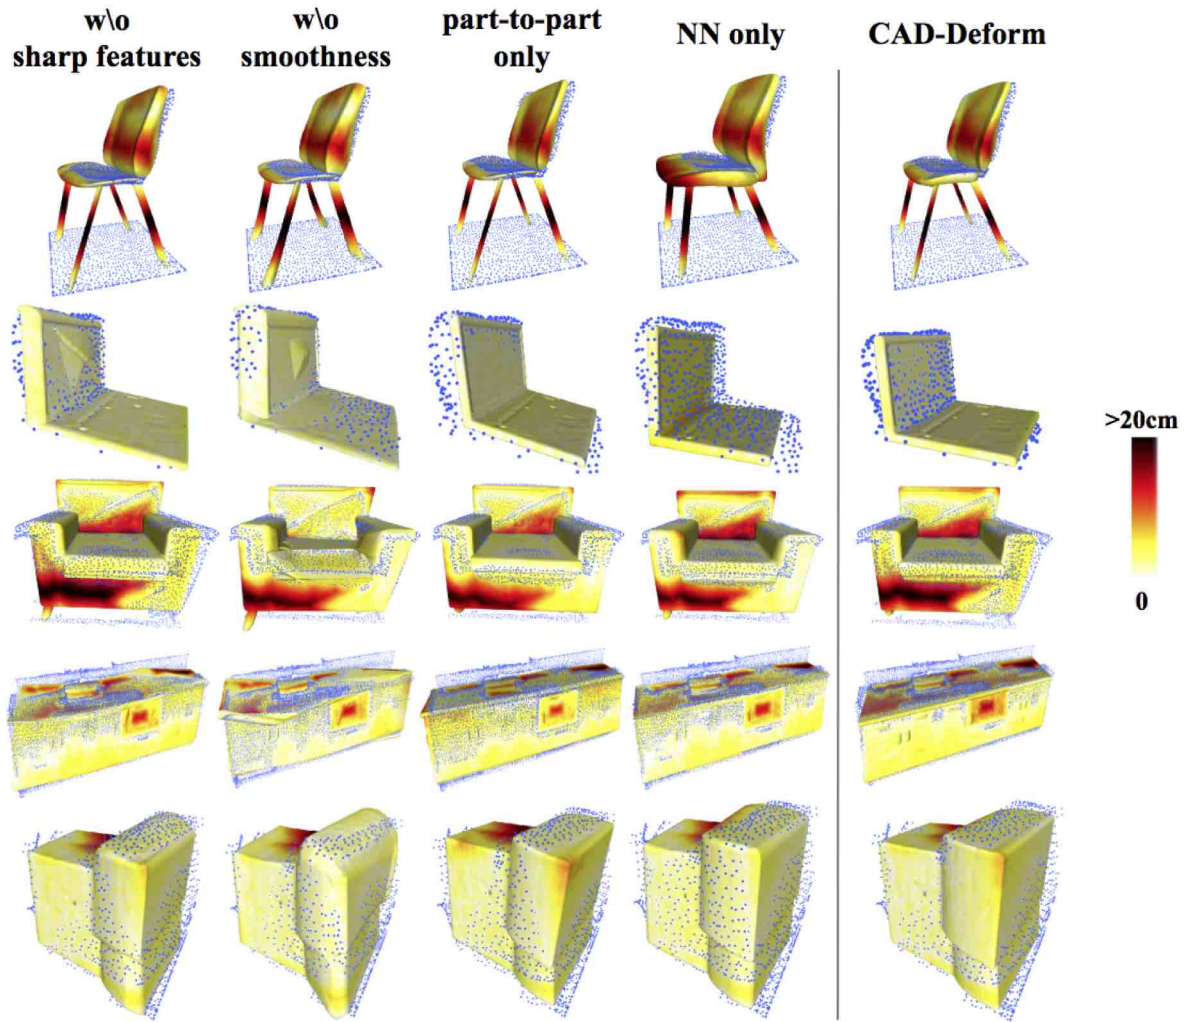


Figure 9: Qualitative results of ablation study using our deformation framework, with mesh coloured according to the value of the tMMD measure.




Figure 10: Qualitative shape translation results, interpolating between the original mesh (left) and the target mesh (right).

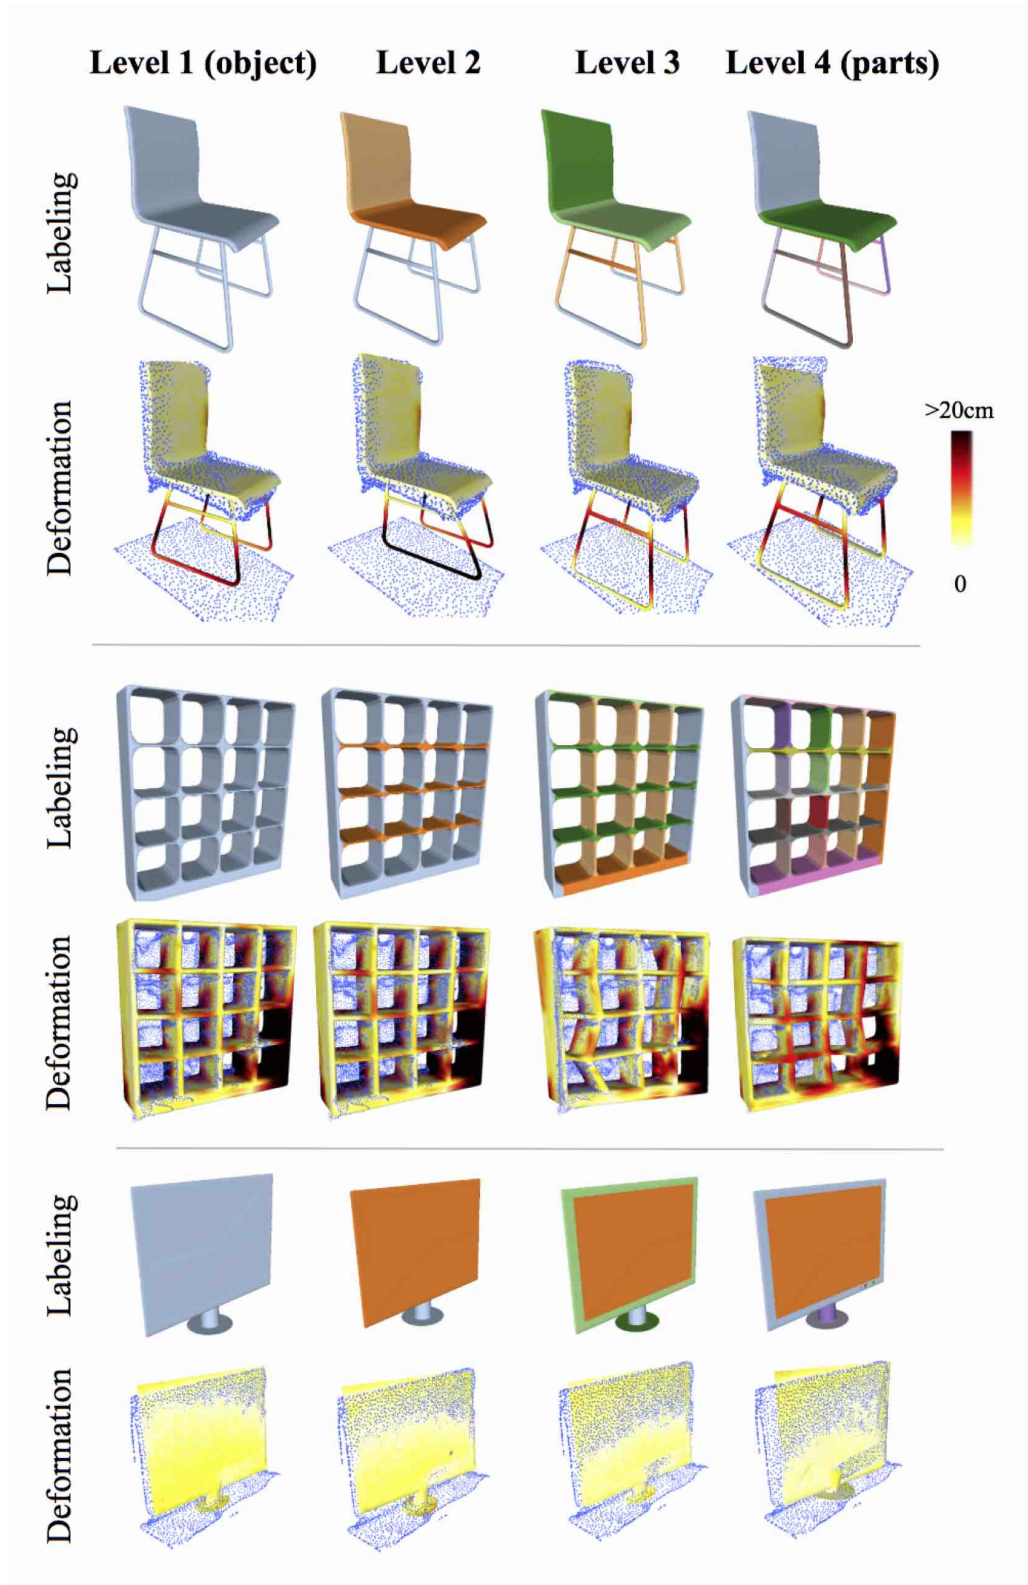


Figure 11: Deformation performance depending on different level of labelling from the PartNet dataset [14]. Deformed mesh surfaces are colored according to the value of tMMD measure, with darker values corresponding to the larger distance values.

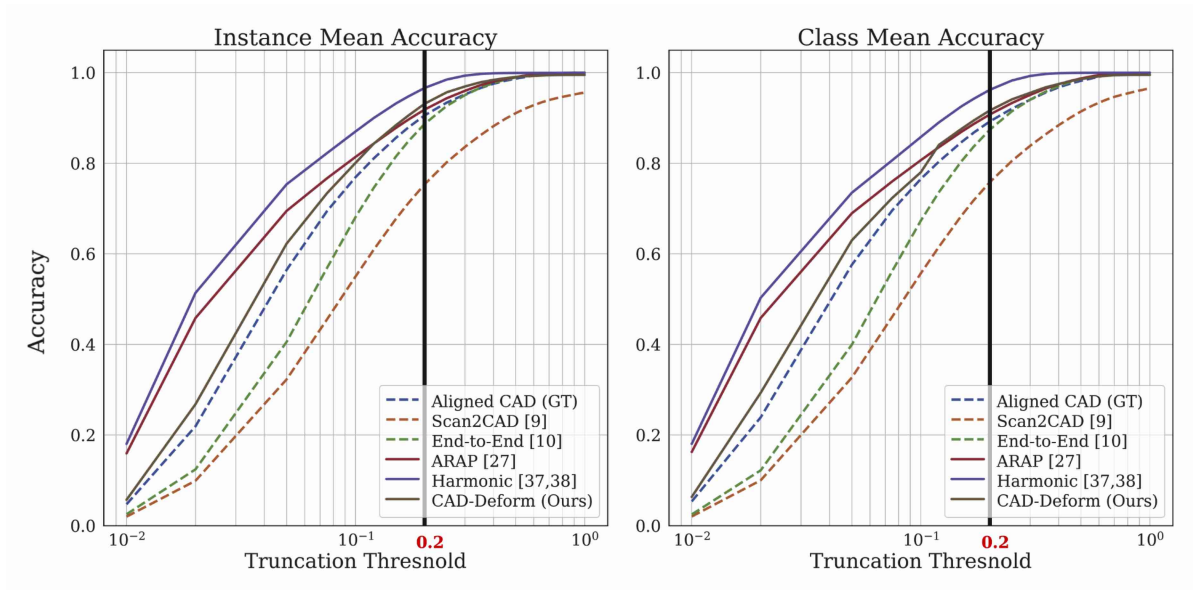


Figure 12: Fitting Accuracy vs. varying τ threshold for the distance between mesh vertices and close scan points.

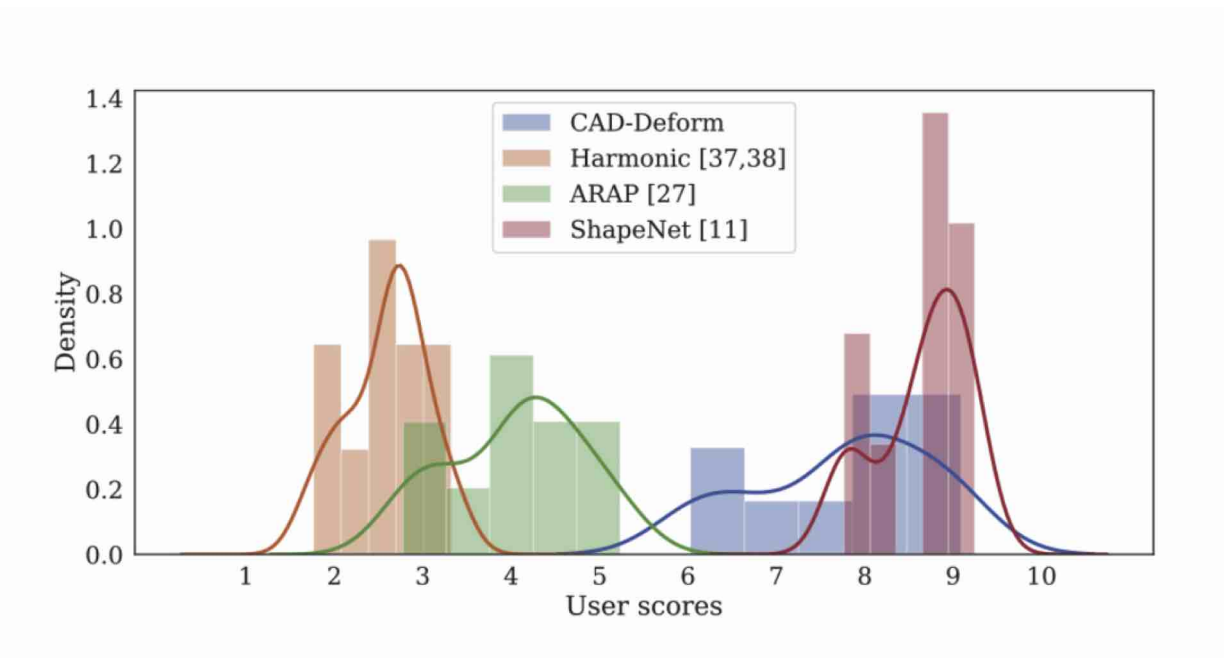


Figure 13: Distribution of user scores averaged by ten shapes from original ShapeNet [11], meshes deformed with ARAP [27], Harmonic [34, 35] and CAD-Deform.