Fit4CAD: A point cloud benchmark for fitting simple geometric primitives in CAD objects

Chiara Romanengo[1,*], Andrea Raffo[1,2], Yifan Qie[3], Nabil Anwer[3], and Bianca Falcidieno[1]

[1]*Istituto di Matematica Applicata e Tecnologie Informatiche "E. Magenes", Consiglio Nazionale delle Ricerche, Via de Marini 6, 16149 Genova, Italy.*
[2]*Department of Mathematics, University of Oslo, Moltke Moes vei 35, 0851 Oslo, Norway.*
[3]*Automated Production Research Laboratory (LURPA), ENS Paris-Saclay, Université Paris-Saclay, 91190 Gif-sur-Yvette, France.*

Abstract

We propose Fit4CAD, a benchmark for the evaluation and comparison of methods for fitting simple geometric primitives in point clouds representing CAD objects. This benchmark is meant to help both method developers and those who want to identify the best performing tools. The Fit4CAD dataset is composed by 225 high quality point clouds, each of which has been obtained by sampling a CAD object. The way these elements were created by using existing platforms and datasets makes the benchmark easily expandable. The dataset is already split into a training set and a test set. To assess performance and accuracy of the different primitive fitting methods, various measures are defined. To demonstrate the effective use of Fit4CAD, we have tested it on two methods belonging to two different categories of approaches to the primitive fitting problem: a clustering method based on a primitive growing framework and a parametric method based on the Hough transform.

Keywords: Benchmarking, Geometric primitive fitting, CAD objects, Quality measures .

1 Introduction

3D CAD models are among the most common medium to convey dimensional and geometric information on designed objects or components. However, often the CAD model of an object is not available, it does not even exist, or no longer corresponds to the real geometry of the manufactured object itself. One strategy for retrieving a digital model of an object when not accessible is to acquire 3D data directly on the object and use it to create a digital representation. The reconstruction of digital models starting from the measured data is a process, commonly called Reverse Engineering (RE), aiming at reconstructing 3D mathematical surfaces and geometric features that represent the geometry of real parts. Many methods have been proposed to solve this problem; as a reference we cite a recent survey that groups a large part of the approaches presented so far [1].

Given the large number of methods proposed, it becomes important to be able to evaluate their performance by creating standard datasets with a ground truth and a "quality label", thus paving the road for a fair evaluation of the existing technologies and the identification of open research directions not only in reverse engineering but also in shape retrieval, understanding, compression, etc., taking inspiration from other approaches proposed for generic classes of objects, (e.g. [2, 3, 4]).

Here we propose Fit4CAD, a benchmark of point clouds representing CAD objects aimed at evaluating methods for detecting simple (polynomial) geometric primitives (i.e., plane, cylinder, cone, sphere, and torus) in 3D point clouds; by polynomial primitive, we here mean a surface that has an algebraic implicit representation, i.e., it can be defined as the zero set of a polynomial. The dataset consists of 225 high quality point clouds, each of which has been obtained by sampling a

---

*Corresponding author

CAD object. Each point cloud is equipped of a ground-truth segmentation and, for each primitive, we provide both implicit and parametric forms. The way these elements were created by using existing platforms and datasets makes the benchmark easily expandable. Fit4CAD is designed to be used also by machine learning methods: in fact, the dataset comes in the form of a training set and a test set.

We provide a number of performance measures able to evaluate both the quality of the fitting segments, in term of points correctly recognized as belonging to a primitive, and the quality of the primitive approximation, evaluating the distance between the primitive detected and the ideal one.

Fit4CAD satisfies a certain number of necessary requirements, such as the relevance and representativeness of the elements in the CAD context, the richness and the completeness of the information associated with each primitive, thus enabling fair comparisons for a wide range of geometric primitives recognition algorithms.

The proposed benchmark has been exploited to evaluate and compare two methods of geometric primitive fitting, belonging to two different categories of approaches: a clustering method based on a primitive growing framework and a parametric method based on the Hough transform. By providing an explicit representation of the equations of the primitives, for the second method we are also able to evaluate measures related to the accuracy of the primitives found.

The rest of the paper is organized as follows. Section 2 examines previous work related to our topic. Section 3 describes the characteristics of the benchmark: dataset, ground truth, and the performance and accuracy measures chosen to evaluate the identification of primitives. Section 4 describes the tests carried out on two methods of recognition and fitting of geometric primitives. Some concluding remarks end the paper.

2   Prior work

Benchmarking involves sharing of resources, metrics, data and so on, so that the common goals of knowledge creation and furthering the state of the art can be achieved. The creation of standard datasets reduces the amount of work necessary for single researchers to assess the quality of their techniques and compare them with other research groups. The steadily rising participation to contests and open challenges shows the interest and the need for benchmarks (e.g., TreCVID [5]), competitive contests (e.g., on Kaggle.com [6] or the 3D Shape Retrieval Contest (SHREC) [7]) and, more in general, for code sharing (e.g., Graphics Replicability Stamp Initiative[1]). So far, benchmarks for 3D object segmentation [8, 9] have mainly considered generic classes of objects and, therefore, the methods were evaluated for their general-purpose segmentation rather than on CAD objects and their capability of recognizing geometric primivites.

Among the datasets containing general 3D shapes (e.g., toys, mechanisms, jewelry) in the form of triangle meshes it is worth mentioning [10, 11, 12], even if these methods were specifically designed for different goals, e.g. 3D printing, computer vision applications and computer graphics applications, respectively.

The most relevant dataset for our work is the ABC dataset [13]; here, the authors present a massive dataset (over one million models), specifically developed to train data-driven algorithms for geometric deep learning. Models are defined by parametric surfaces, possibly accompanied with the information related to the decomposition into patches, sharp feature annotations, and analytic differential properties. The models were created by using the interface available on the online infrastructure Onshape[2]. All models are stored as triangle meshes, while the associated files (annotations, features, etc.) are not available for all models: more precisely, the file may lack the true list of primitives, or their parametric/implicit representations; this is not a big issue for that specific dataset, as primitive extraction is not their purpose. Being specifically designed for data-driven methods, the models are stored with different resolutions (i.e., different samplings of the same parametric model) or slight variations. Moreover, models in [13] do not present any kind of data

---

[1]http://www.replicabilitystamp.org
[2]https://www.onshape.com/en/

perturbation. Lastly, the ABC dataset is not designed for point cloud segmentation, and does not present any specific quality measures for comparing methods.

3 The benchmark

We here introduce our benchmark, geared towards the following desirable properties:

- *Dataset richness and representativeness.* The first and foremost requirement for a thorough evaluation is the availability of a data set characterized by a good sampling of varied shapes: each family of simple geometric primitives (i.e., plane, cylinder, cone, sphere, and torus) should appear in a sufficient number of point clouds. In addition, we consider different point cloud densities as well as data integrity (with/without missing data). Each CAD object is used to generate one and only one point cloud.

- *Ease of expansion.* An ideal benchmark should be able to develop over time, in order to test new paradigms and face new challenges; this requires the capability to generate new data in an easy and efficient way. To satisfy this basic requirement, Section 3.1 outlines a general pipeline that can be used for data generation.

- *Availability of both implicit and parametric representations.* Modern CAD systems are based on two complementary representations for surfaces, according to the manipulation they are involved in: implicit and parametric representations. Parametrized surfaces are best suited for point generation, while implicit representations allow to check whether a query point lies or not on the surface in a more convenient way. Having both representations makes it possible to answer a wide range of questions (e.g., intersection problems).

- *Completeness of the documentation.* All models are equipped with all and the same information. For each of them, this includes: the files with primitive segments, implicit and parametric primitive representations; a preset split of the dataset into training set and test set. Further details are provided in Section 3.2.

- *Variety of performance indicators and accuracy measures.* To evaluate and compare methods, it is of vital importance to select measures that highlight strengths and weaknesses. In our case, the problem is twofold: on the one hand, we want to quantify the capability to produce precise segmentations into simple geometric primitives (by performance indicators); on the other hand, we also aim at measuring the fitting accuracy when it comes to implicit and parametric representation of the same shapes (by accuracy measures). Section 3.3 describes the measures selected to evaluate the detection of simple primitives in CAD objects point clouds.

3.1 Dataset

At present, the dataset contains 225 individual high quality point clouds, each of which has been obtained by sampling a CAD object. The dataset is already split into two subsets: a training set, counting 190 point clouds, and a test set, containing the remaining 35 point clouds. Figure 1 shows the distribution of surface types for both training and test sets.

The dataset generation process, in the most general form, has been carried out by the following three steps:

1. *Model creation.* We created part of the models, by using the publicly available interface hosted by Onshape, while the remaining part was collected from the ABC dataset [13], which was derived, in turn, from the Onshape public collection. Models gathered from the ABC dataset have been filtered by manually correcting the parts presenting minimum flaws and rejecting low quality models, in order to avoid rare yet bothersome imperfections, such as overlapping or repeating patches. Some examples of CAD objects from Onshape are displayed in Figure 2.
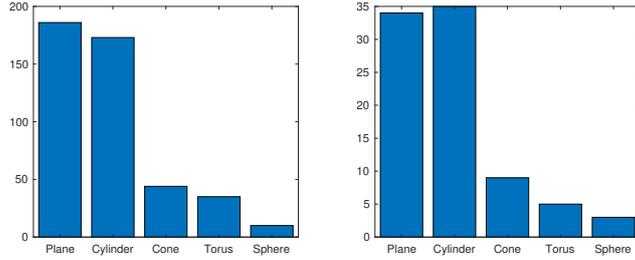
Figure 1: Surface type distribution. The two bar charts show the distributions for the training set (left) and the test set (right).

2. *Parametric and implicit representations.* The generation of B-rep models was crucial to extract the parametric representation behind each geometric primitive; in our case, the parametric representations for each patch have been obtained by processing the STEP files produced by Onshape in GMSH [14]; nevertheless, we emphasize that other software could be considered too (e.g., [15]). Several methods to compute the implicit representation from a parametric form are nowadays available. We here consider the numerical approach known as *approximate implicitization*, introduced in [16] and further delevoped in [17]. One of the advantages of this approach is that it provides exact implicit representations when the exact total degree is selected; we remind that a bivariate polynomial has total degree n if all monomials $x^i y^j$ are such $i + j \leq n$, and there exists at least one monomial $x^i y^j$ such that $i + j = n$.

3. *Point cloud extraction.* CAD objects are sampled at different densities, and optionally manually postprocessed by using CloudCompare[3] to simulate missing data. To give an example, Figure 3(a) shows a model from Onshape, which is then sampled and postprocessed in 3(b-c).
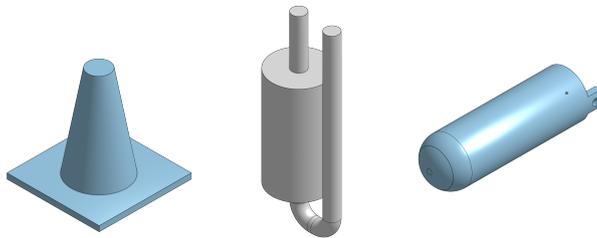
Figure 2: Example of models obtained using Onshape.

(a)   (b)   (c)
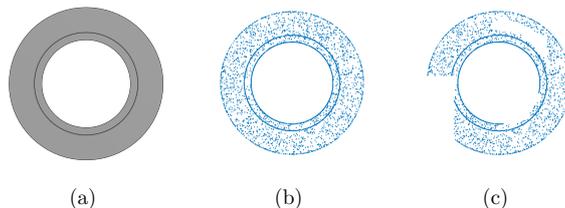
Figure 3: Example of point cloud creation. The initial object in (a) is sampled at a chosen density (b) and then perturbed by simulating missing data (c).

---

[3]CloudCompare (version 2.10.2), http://www.cloudcompare.org/
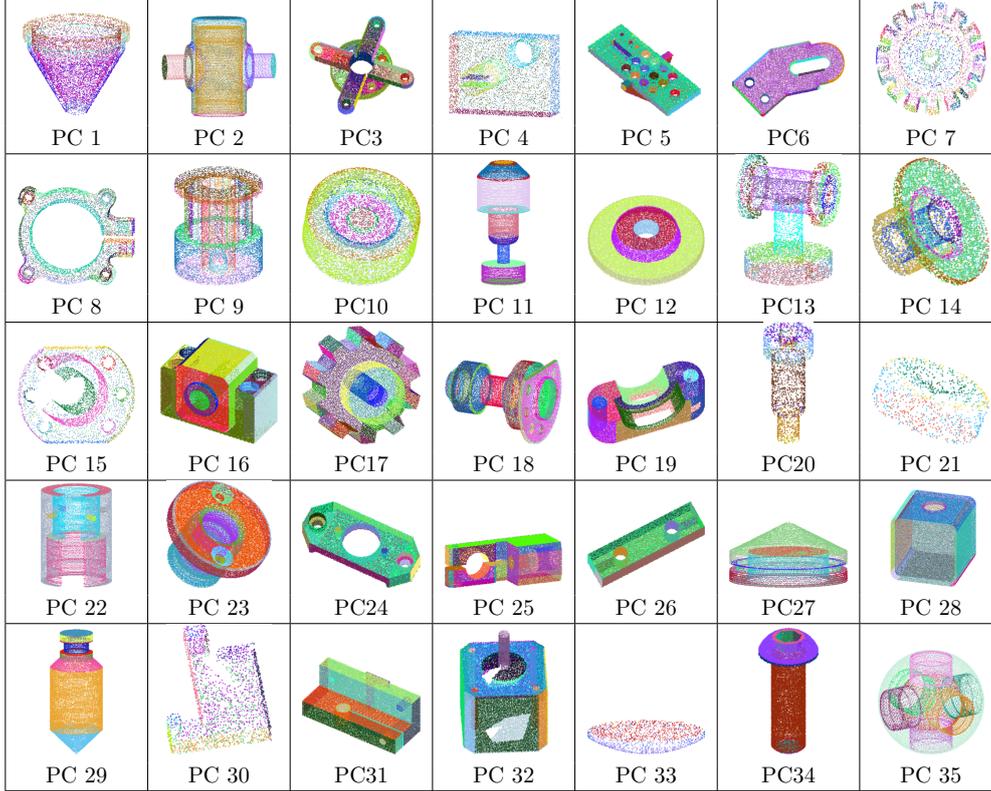
Figure 4: The 35 point clouds used as a test set. Different colors represent different primitives, as stored in the CAD models, i.e., our ground truth.

3.2   Ground truth

Each model in the ground truth comes in the form of four TXT files. We here provide a description of each file content for the i-th point cloud.

PCi lists the three-dimensional points forming the point cloud to be segmented.

PCi_primitives contains the list of true primitives. For each primitive, a list of indices is provided; each index corresponds to a point in "*PCi*", with respect to the ordering there introduced. For example,

```
Primitive6:=[4 9 184 185 186 187 188 189 190 191 192]
```

means that the sixth primitive contains points number 4, 9, 184, 185, 186, 187, 188, 189, 190, 191 and 192 (where the ordering is the one in the corresponding "*PCi*").

PCi_parametric provides, for each primitive in "*PCi_primitives*" corresponding to a plane, a cylinder, a cone, a sphere or a torus, its parametric representation. To give an example,

```
Primitive6:=[primitive type, v]
```

where **v** is the vector that contains the parameters of the parametric representation (see A.1 for further details on the considered ordering).

PCi_implicit provides, for each primitive in "*PCi_primitives*" corresponding to a plane, a cylinder, a cone, a sphere or a torus, its implicit representation. For example,

```
Primitive6:=[primitive type, w]
```

5

where **w** is the vector that contains the coefficients of the implicit representation (see A.2 for further details on the considered ordering).

The points that do not correspond to any of the simple primitives mentioned above (i.e., plane, cylinder, cone, sphere or torus) are classified as unsegmented and not explicitly reported in files *PCi_primitives*, *PCi_parametric* and *PCi_implicit*; in the original model, these points usually originate from B-spline surfaces. We intentionally decided to insert some models with non-simple geometric primitives to check whether a candidate method can avoid misclassification.

3.3   Quality indicators

To evaluate the detection of simple primitives in CAD objects, we have proposed quality measures selected from [18, 19] with particular care on what concerns their performance and approximation accuracy.

3.3.1   Performance measures of the point classification

Any primitive in a model is identified by the list of points belonging to it or, equivalently, by the list of points that do not belong to it. The problem of primitive detection can therefore be easily written in terms of binary classification tasks, one per primitive in the ground truth.

Let $\mathscr{P}_B$ be a a set of points in the benchmark point cloud corresponding to a specific primitive, and let $\mathscr{P}_S$ be the primitive in the segmentation to assess that most overlap with $\mathscr{P}_B$. We can define the following quantities:

- *True positives*, TP: the number of points shared by $\mathscr{P}_B$ and $\mathscr{P}_S$.

- *False positives*, FP: the number of points in $\mathscr{P}_S$ that do not belong to $\mathscr{P}_B$.

- *False negatives*, FN: the number of points in $\mathscr{P}_B$ that do not belong to $\mathscr{P}_S$.

- *True negatives*, TN: the number of points that do not belong to either $\mathscr{P}_B$ nor $\mathscr{P}_S$.

Based on these four quantities, we consider the following measures:

- *Sensitivity*, also called *true positive rate*, measures the proportion of positives which are correctly identified, i.e.,
$$\mathrm{TPR} := \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}.$$

  Specificity, or *true negative rate*, measures the proportion of true negatives that are correctly identified as such, i.e.,
$$\mathrm{TNR} := \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FP}}.$$

- *Positive predictive value* is defined as the proportion of predicted positives which are actual positives, i.e.,
$$\mathrm{PPV} := \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}.$$

  Similarly, *negative predictive value* is given by
$$\mathrm{NPV} := \frac{\mathrm{TN}}{\mathrm{TN} + \mathrm{FN}}.$$

- *Accuracy* is the ratio of correct predictions to total predictions made, i.e.,
$$\mathrm{ACC} := \frac{\mathrm{TP} + \mathrm{TN}}{\mathrm{TP} + \mathrm{TN} + \mathrm{FP} + \mathrm{FN}}.$$

- *Sørensen-Dice index.* It is given by

$$\mathrm{DSC} := \frac{2|\mathscr{P}_B \cap \mathscr{P}_S|}{|\mathscr{P}_B| + |\mathscr{P}_S|}.$$

In case of binary classification, it is shown to be equivalent to

$$\mathrm{DSC} := \frac{2\mathrm{TP}}{2\mathrm{TP} + \mathrm{FP} + \mathrm{FN}},$$

which is often referred to as F_1 *score.*

For more details, we refer the reader to [18].

3.3.2 Approximation accuracy

To measure the recognition accuracy of a specific primitive, we use the parametric and the implicit representations provided in "*PCi_implicit*" and "*PCi_parametric*". Exploiting the notation provided before, let us consider a primitive $\mathscr{P}_S$ to be evaluated, and let $\mathcal{S}$ be the surface described by the corresponding parametric representation. When it comes to the parametric representation, we use the following two measures to evaluate the approximation accuracy of primitive $\mathscr{P}_S$:

- *Mean Fitting Error* (MFE):

$$\mathrm{MFE}(\mathscr{P}_S, \mathcal{S}) := \frac{1}{|\mathscr{P}_S|} \sum_{\mathbf{x} \in \mathscr{P}_S} d(\mathbf{x}, \mathcal{S})/l, \tag{1}$$

where d is the Euclidean distance, and l is the diagonal of the minimum bounding box containing $\mathscr{P}_S$.

- *Directed Hausdorff distance*:

$$d_{\mathrm{dHaus}}(\mathscr{P}_S, \mathcal{S}) = \max_{\mathbf{x} \in \mathscr{P}_S} \min_{\mathbf{y} \in \mathcal{S}} d(\mathbf{x}, \mathbf{y}),$$

with d the Euclidean distance. To make the measure independent from the primitive size, we normalize it with respect to the diagonal l of the minimum bounding box containing $\mathscr{P}_S$.

The fitting accuracy for the implicit representation is evaluated by the following measure:

- *Coefficient distance*:

$$d_1(\mathbf{v}, \mathbf{v}') = \|\mathbf{v} - \mathbf{v}'\|_1$$

where $\mathbf{v}$ and $\mathbf{v}'$ are the coefficient vectors for the implicit representations of the primitives $\mathscr{P}_S$ and $\mathscr{P}_B$, respectively, and where $\|\cdot\|_1$ is the well-known ℓ^1 norm. In order to make this measure consistent, we assume the coefficient vectors to be normalized, and the first nonzero entry to be positive (where the ordering is the one provided in A.2).

We refer the reader to [19] for further details.

4 Test of the benchmark on two methods

The proposed benchmark has been used to evaluate and compare two methods dealing with primitive fitting. As guiding examples of how the benchmark works, we have selected two methods that are both available and representative of two classes of methods according to the taxonomy defined in [1]: a clustering method based on a primitive growing framework (Section 4.1) and a parametric method based on the Hough transform [20] (Section 4.2). Both methods can be evaluated according to the measures described in Section 3.3.1, as they explicitly provide the list of the points that form any primitive; on the other hand, the approximation accuracy can be assessed only for methods that can provide parametric or implicit representations, in our examples the Hough-based fitting.

4.1 PG: a discrete curvature-based method for point cloud segmentation

As a first approach, we present a curvature-based method based on a primitive growing framework, on the basis of the method proposed in [21] for triangle meshes; for the sake of brevity, we will often use the acronym PG as a shorthand for this method, where PG stands for "Primitive Growing".

The method consists of two main steps: an initial region partitioning process based on high curvature detection and, then, a region refinement process based on slippage analysis, as summarized in Figure 5. These two steps run as follows:

- *Initial region partition.* Points that identify sharp edges are characterized via a point attribute called *surface variation*, as introduced in [22]. Given a point $\mathbf{p}$ and n neighbouring points, its surface variation is defined as

$$\sigma_n(\mathbf{p}) := \frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3},$$

  where $\lambda_1 \leq \lambda_2 \leq \lambda_3$ are the eigenvalues of the covariance matrix for the sample point $\mathbf{p}$ and its n neighbouring points; note that λ_l measures the variability of the neighborhood of n points along the direction of the corresponding eigenvector. In our experiments, n is set to 15 as it yields good results when considering the training set. Points on sharp edges are characterized by a high surface variation. These points are here selected by analysing a histogram of the surface variation values, by means of a threshold η: all points having surface variation above η are labelled as sharpe edge points; for example, setting $\eta = 0.8$ means that the points whose surface variation is higher than 80% (top 20%) are considered to belong to a sharp edge. In our implementation, the threshold η is user-defined and taken in the interval $[0.65, 0.95]$. Once sharp edges have been identified, a region growing approach is applied to compute a first coarse pre-segmentation, along the lines of what detailed in [22]: starting from a random seed point, its nearest neighbors are progressively located; those points which does not belong to sharp edges will be labelled and used as new seed points, until all neighboring points are labelled.

- *Region refinement.* According to the ISO GPS invariance class [23], "ideal" features can be categorized into seven invariance classes: planar, cylindrical, helical, spherical, revolute, prismatic, and complex. The seven invariance classes are here captured by local slippage analysis [24]. This step aims at decomposing any coarse segment S from the previous step into simpler geometric parts. Given a point set P of n points from S, the slippable motions of P are found as the motion vector $[\mathbf{r}, \mathbf{t}]$ that, when applied to P, minimizes the motion along the normal direction at each point

$$\min_{[\mathbf{r},\mathbf{t}]} \sum_{i=1}^{n} ((\mathbf{r} \times \mathbf{x}_i + \mathbf{t}) \cdot \mathbf{n}_i)^2, \tag{2}$$

  where: $\mathbf{r} = (r_x, r_y, r_z)$ is a rotation vector around x, y, and z; $\mathbf{t} = (t_x, t_y, t_z)$ is a translational vector; $\mathbf{p}_i \in P$ are the n samples, and $\mathbf{n}_i$ are their respective normals. Equation 2 is a least-square problem which can be reduced to the linear system based on the covariance matrix of the second partial derivatives of the function in 2 with respect to the rotation and translation parameters, see [24] for further details. Slippage analysis permits the detection of 3-, 2-, 1- and 0-slippable motions, see Table 1. According to the primitives defined in the benchmark, segments identified as prismatic, revolute and complex could be undersegmented; to address this problem, the RANSAC method introduced in [25] is applied; an example of prismatic segment requiring further processing is shown in Figure 6. Finally, points on sharp edges are assigned to the closest primitives.

4.1.1 Computational complexity

The characterization of the sharp edges according to surface variation [22] is done by considering the k-nearest neighbour points with $\mathcal{O}(n \log n)$ operations, where n represents the number of points. The growing algorithm for grouping points inside boundaries costs $\mathcal{O}(n)$ operation while the classification
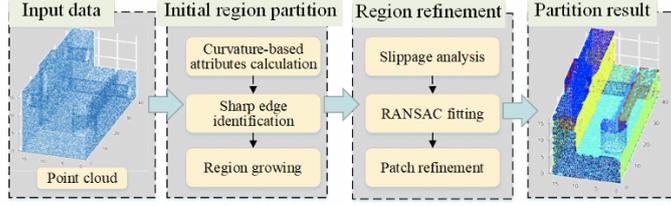
Figure 5: The framework of the curvature-based surface partitioning method.

ISO GPS Invariance	Slippage	geometric primitives
planar	3	plane
spherical	3	sphere
cylindrical	2	cylinder
helical	2	-
prismatic	1	undersegmented
revolute	1	cone/torus
complex	0	undersegmented

Table 1: Relation between ISO GPS invariance class [23] and the simple geometric primitives in this benchmark. Note that, in this terminology, prismatic/complex could include include planar, spherical or cylindrical primitives.
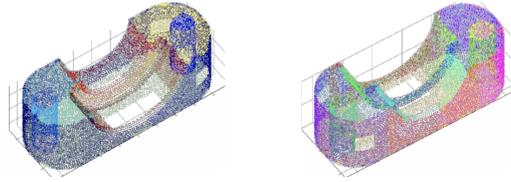
Figure 6: Example of undersegmented invariance class. On the left, the external blue primitive is classified as prismatic. On the right, the same primitive is split into 2 planes and 2 half cylinders.

of each point set via slippage analysis is $\mathcal{O}(m)$, where m is the number of points in one surface portion [24]. The further RANSAC based segmentation in case of point sets with low slippage values (such as revolute, prismatic and complex primitives) is $\mathcal{O}(m)$ [25], where $m < n$ in the most common scenario.

4.2   HT: Simple primitive fitting based on Hough transform
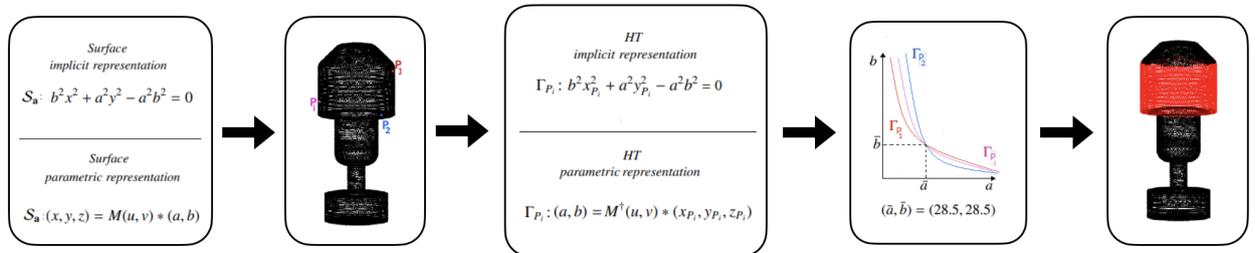
Figure 7: The HT-based paradigm: a visual illustration of how a cylinder representation S_a is converted by the Hough transform into $|P_i|$ hypersurfaces Γ_{P_i}; then the intersection of the hypersurfaces Γ_{P_i} identifies the parameters $(\bar{a}, \bar{b})$ that correspond to the red cylinder in the right.

In this section, we consider a method to segment and to fit a point cloud PC with surface primitives using the Hough Transform (HT) technique. The general HT-framework deals with the

9

problem of finding a surface $\mathcal{S}_{\bar{\mathbf{a}}}$ – within a family $\mathcal{F} = \{\mathcal{S}_{\mathbf{a}}\}$ of surfaces dependent on a set of parameters $\mathbf{a} = (a_1, ..., a_n)$ – that best approximates a particular shape. The common strategy to identify the solution (or a solution) consists in a procedure whereby each point in PC votes a n-uple $\mathbf{a}$ in the parameter space; the most voted n-uple $\bar{\mathbf{a}}$ corresponds to the most representative surface $\mathcal{S}_{\bar{\mathbf{a}}}$ for a dense subset of PC. Figure 7 illustrates how the Hough transform converts the problem of fitting points on a primitive into the problem of fitting the parameters of a family of primitives into points.

This method is based on the theory related to the extension of the Hough transform to general algebraic objects [26]. This theory is very broad and can be used for many types of primitives, for instance in [27] is used for fitting point sets with ellipsoids and can deal with non-simple primitives, such as helical surfaces. The families of primitives included in this benchmark are planes, cylinders, spheres, cones and tori. Once a family of primitives $\mathcal{F}$ is selected, the main steps can be summarized as follows:

- *Inizialization and estimation of the accumulator function.* Once the family $\mathcal{F}$ is chosen, a region T of the parameter space is selected exploiting the knowledge of the geometric characteristics of $\mathcal{F}$ (e.g., bounding box). Then, it is discretized into cells, which are uniquely identified by the coordinates of their centre. This space is associated with an accumulator function $\mathcal{H}$, discretized as a matrix. Its entries are in a one-to-one correspondence with the cells of T. An entry of $\mathcal{H}$ is increased by 1 each time the HT of a point P, Γ_P, intersects the corresponding cell.

- *Selection of potential fitting primitives.* In the case the input point cloud is composed of different primitives, the peaks of $\mathcal{H}$ identify the potential primitives $\mathcal{S}_{\bar{\mathbf{a}}_{\mathbf{i}}}$ that might fit different parts $\mathcal{X}_i \subseteq PC$. Then, the cells corresponding to the peak values of the accumulator function $\mathcal{H}$ are identified by studying its *topological persistence* (see [28]). In our implementation, the peaks that correspond to primitives are automatically recognised by keeping the local maxima with a persistence higher than 10% of the maximum value of $\mathcal{H}$, using the algorithm for persistent maxima proposed in [29]. The coordinates of the cell centres of the maxima or the peaks of the accumulator function correspond to the parameters of potentially recognised surface primitives.

Since it can happen that more types of primitives fit the same dense subset $\mathcal{X}_i$ (or a part of it), the *Mean Fitting Error* (see Equation 1) is used to evaluate the approximation accuracy of each primitive. Then, if $\mathcal{S}_{\bar{\mathbf{a}}_{\mathbf{i},\mathbf{1}}}$ and $\mathcal{S}_{\bar{\mathbf{a}}_{\mathbf{i},\mathbf{2}}}$ are two candidate primitives, the fitting errors $\mathrm{MFE}(\mathcal{X}_i, \mathcal{S}_{\bar{\mathbf{a}}_{\mathbf{i},\mathbf{1}}})$ and $\mathrm{MFE}(\mathcal{X}_i, \mathcal{S}_{\bar{\mathbf{a}}_{\mathbf{i},\mathbf{2}}})$ between each primitive and $\mathcal{X}_i$ are calculated; the primitive having lowest error is kept. The final result is the partitioning of the input point cloud PC into several subsets in such a way that points of the same segment are well approximated by the same primitive. Figure 8 summarizes the HT framework. In particular, Figure 8(b) shows an example of accumulator matrix referred to the recognition of the cylinders, while the four cylinders corresponding to the four peaks are highlighted on the original point cloud in Figure 8(c). Finally, Figure 8(d) exhibits the resulting segmentation.

4.3 Evaluation

We here analyse the performance of the methods outlined in Sections 4.1 and 4.2, with the purpose of showing how the benchmark works. Firstly, we compare the quality of the segments/primitives found against a ground-truth; the measures involved do not require an explicit representation of the primitive equation, and thus can be applied to both methods. Secondly, we consider the accuracy of the parametric/implicit representations: in our case, this comes down to the analysis of the method introduced in Section 4.2.

Performance measures of the point classification

Figures 11 and 12 in B provide the segmentation results obtained by the primitive growing (PG) and Hough transform (HT) based approaches, colored as follows: given a model and an approach,
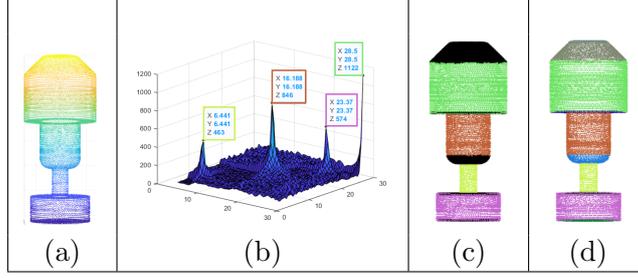
Figure 8: HT framework. In (a), an example of input point cloud; in (b), the accumulator function associated with the search for cylinders and the peaks found by the method for persistent maxima; in (c), the four cylinders corresponding to the four peaks (black stands for unclassified points). The final outcome, after searching for all simple geometric primitives, is shown in (d).

for each primitive in the benchmark we find the most overlapping segmented primitive; misclassified points are colored in black, while correct matches follows the $1 - 1$ primitive-color correspondence from Figure 4.

Table 2 summarizes the performances of the two methods over all the test set models. Each row correspond to a model; for each model, the table provides information on the number of true and predicted primitives, as well as the accuracy measures introduced in Section 3.3.1. For each metric, two columns are considered, respectively referring to the PG- and the HT-based approaches.

	# points	# true primitives	# predicted primitives		DSC		PPV		TPR		TNR		NPV		ACC	
			PG	HT	PG	HT	PG	HT	PG	HT	PG	HT	PG	HT	PG	HT
PC 1	7,500	8	12	8	0.370	0.987	0.693	0.989	0.332	0.985	0.968	0.998	0.696	0.995	0.706	0.995
PC 2	20,621	17	13	16	0.656	0.921	0.795	0.960	0.632	0.937	0.994	0.999	0.947	0.998	0.944	0.998
PC 3	9,723	35	36	35	0.843	0.896	0.778	0.828	0.943	0.986	0.994	0.996	0.998	1.000	0.992	0.996
PC 4	10,000	15	12	12	0.736	0.840	0.831	0.981	0.729	0.800	0.993	0.999	0.964	0.972	0.960	0.973
PC 5	20,000	37	32	34	0.480	0.839	0.665	0.872	0.544	0.901	0.990	0.999	0.925	0.978	0.918	0.978
PC 6	9,320	26	13	20	0.465	0.783	0.671	0.857	0.466	0.772	0.994	0.999	0.968	0.997	0.964	0.996
PC 7	5,000	68	41	69	0.465	0.923	0.642	0.871	0.429	0.994	0.993	0.999	0.978	1.000	0.972	0.999
PC 8	7,500	32	28	30	0.390	0.890	0.528	0.896	0.438	0.889	0.984	0.997	0.930	0.998	0.916	0.996
PC 9	17,000	104	43	40	0.403	0.568	0.607	0.848	0.335	0.456	0.998	0.999	0.993	0.996	0.991	0.995
PC 10	10,000	10	7	10	0.627	0.919	0.793	0.926	0.555	0.938	0.988	0.995	0.949	0.996	0.943	0.993
PC 11	13,201	13	9	10	0.705	0.805	0.874	0.959	0.677	0.770	0.993	0.998	0.937	0.973	0.934	0.973
PC 12	12,327	6	6	6	0.894	0.998	0.852	0.997	0.982	1.000	0.984	0.999	0.994	1.000	0.981	0.999
PC 13	10,000	21	17	16	0.582	0.800	0.792	0.991	0.576	0.761	0.993	1.000	0.951	0.983	0.946	0.984
PC 14	7,500	8	6	8	0.623	0.963	0.836	0.997	0.606	0.933	0.987	0.999	0.878	0.990	0.878	0.991
PC 15	5,000	15	15	14	0.533	0.941	0.615	0.972	0.566	0.933	0.981	0.999	0.951	0.993	0.937	0.992
PC 16	29,641	35	28	31	0.676	0.848	0.771	0.842	0.702	0.891	0.997	0.999	0.986	0.994	0.984	0.993
PC 17	21,137	45	45	45	0.916	0.935	0.859	0.889	0.983	0.996	0.997	0.998	1.000	1.000	0.997	0.998
PC 18	16,406	29	20	29	0.555	0.933	0.847	0.912	0.536	0.978	0.994	0.999	0.898	0.999	0.896	0.998
PC 19	16,740	16	14	11	0.751	0.747	0.869	0.967	0.747	0.681	0.990	0.998	0.971	0.960	0.963	0.960
PC 20	2,500	14	8	14	0.525	0.954	0.789	0.917	0.467	1.000	0.991	0.999	0.916	1.000	0.914	0.999
PC 21	1,000	5	3	5	0.677	0.985	0.891	0.983	0.588	0.988	0.977	0.997	0.834	0.998	0.845	0.996
PC 22	26,093	10	6	10	0.561	0.967	0.751	0.949	0.586	0.987	0.994	1.000	0.824	1.000	0.822	0.999
PC 23	19,088	13	14	12	0.721	0.886	0.699	0.878	0.828	0.930	0.985	0.997	0.987	0.991	0.975	0.989
PC 24	13,767	27	21	27	0.742	0.916	0.795	0.861	0.750	0.999	0.995	0.997	0.992	1.000	0.987	0.998
PC 25	18,331	38	26	35	0.677	0.873	0.841	0.862	0.665	0.921	0.997	0.998	0.986	0.998	0.984	0.996
PC 26	17,374	14	10	14	0.663	0.975	0.762	0.953	0.607	1.000	0.988	0.999	0.965	1.000	0.956	0.999
PC 27	19,339	9	7	9	0.789	0.994	0.860	0.995	0.751	0.993	0.988	1.000	0.971	0.999	0.963	0.999
PC 28	46,364	21	15	21	0.589	0.983	0.776	0.975	0.551	0.993	0.996	0.999	0.963	0.999	0.960	0.999
PC 29	12,753	9	7	9	0.785	0.991	0.852	1.000	0.758	0.983	0.992	1.000	0.985	0.998	0.980	0.999
PC 30	2,500	13	8	12	0.468	0.873	0.666	0.873	0.421	0.887	0.974	0.995	0.863	0.978	0.855	0.974
PC 31	22,098	22	18	23	0.729	0.869	0.813	0.906	0.758	0.887	0.992	0.995	0.980	0.991	0.972	0.986
PC 32	18,950	22	18	22	0.682	0.972	0.853	0.948	0.749	1.000	0.994	0.998	0.992	1.000	0.987	0.998
PC 33	1,500	3	3	3	0.811	0.978	0.875	0.983	0.819	0.974	0.910	0.983	0.948	0.997	0.901	0.988
PC 34	12,089	14	13	14	0.788	0.939	0.823	0.892	0.806	0.998	0.994	0.998	0.979	0.999	0.976	0.998
PC 35	24,068	8	15	8	0.866	0.975	0.872	0.958	0.960	0.999	0.983	0.992	0.997	1.000	0.984	0.995

Table 2: Number of fitted primitives and classification performance metrics: comparison between the PG-based and the HT-based algorithms.

To ease the analysis, the metrics are studied via boxplots:

- Figure 9 compares the two methods over the whole test set. A first observation of this analysis is that accuracy measures from the HT-approach have generally a lower variability. At a closer

11

look, one can notice that the quartiles, as well as the minimum and maximum, always assume higher values when it comes to the HT-based method; in particular, the second quartile (i.e., the median) is always above 90%. DSC, TPR and PPV are the three accuracy measures that varies the most; this highlights that the two methods have lower performances in identifying the true positives, compared to true negatives. Both methods exhibit outliers in most of the boxplots.

- Robustness to missing data is analysed in Figure 10. The HT-based method turns out to be hardly affected by such perturbation, as the inter-quartile range and the whiskers do not significantly vary; the only noteworthy variation is that of TPR, which points out a slightly decreased capability in correctly identifying positives. A more prominent variation can be noted for the PG-method.

Interestingly enough, both methods rarely suffer from oversegmentation, while it is more likely for them to undersegment. The most dramatic undersegmentation is that of point cloud 9 (i.e., PC 9 in Table 2), where the PG-based and the HT-based methods only manage to detect 43 and 40 primitives, respectively, out of the 104 there expected; this highlights possible issues when the original model has thin or small primitives.
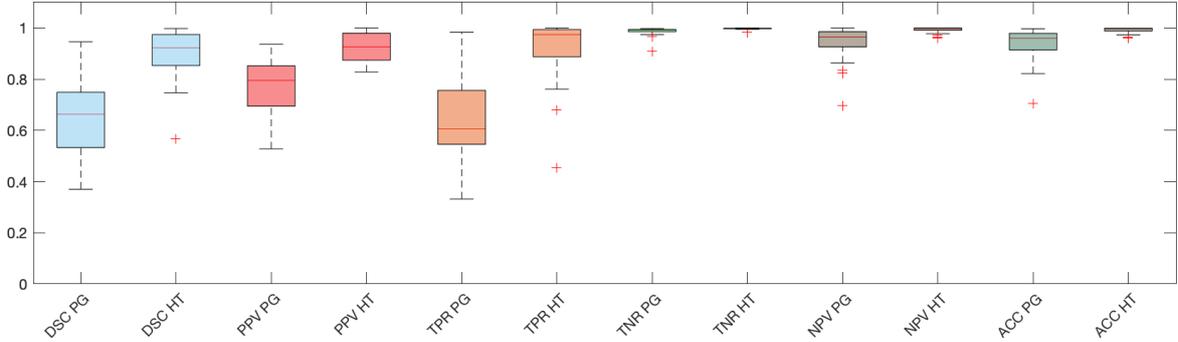
Figure 9: Boxplot for the classification metrics presented in Table 2. All 35 models are here considered.

Approximation accuracy

Table 3 reports the performance of the HT-based method evaluated according to the metrics reported in Section 3.3.2. Each row corresponds to a segmented point cloud from Figure 12; each column represents a different accuracy measure; for each point cloud, each measure has been obtained by averaging over all segments.

- Being the MFE normalized by definition, its value can be interpreted as a percentage. From the numbers provided in the table, we can conclude that the MFE ranges from a minimum of 0.1% to a maximum of 1.0%.

- The directed Hausdorff distance, in its normalized version, ranges from 0.2% to 1.9%. The generally higher values, compared to those from the MFE, can be explained by the Hausdorff's sensitivity to outliers.

- The coefficient distance seems to provide a much more fluid situation. By checking the model corresponding to the highest error, we can conclude that the HT-based method has lower precision when applied to point clouds containing tori.
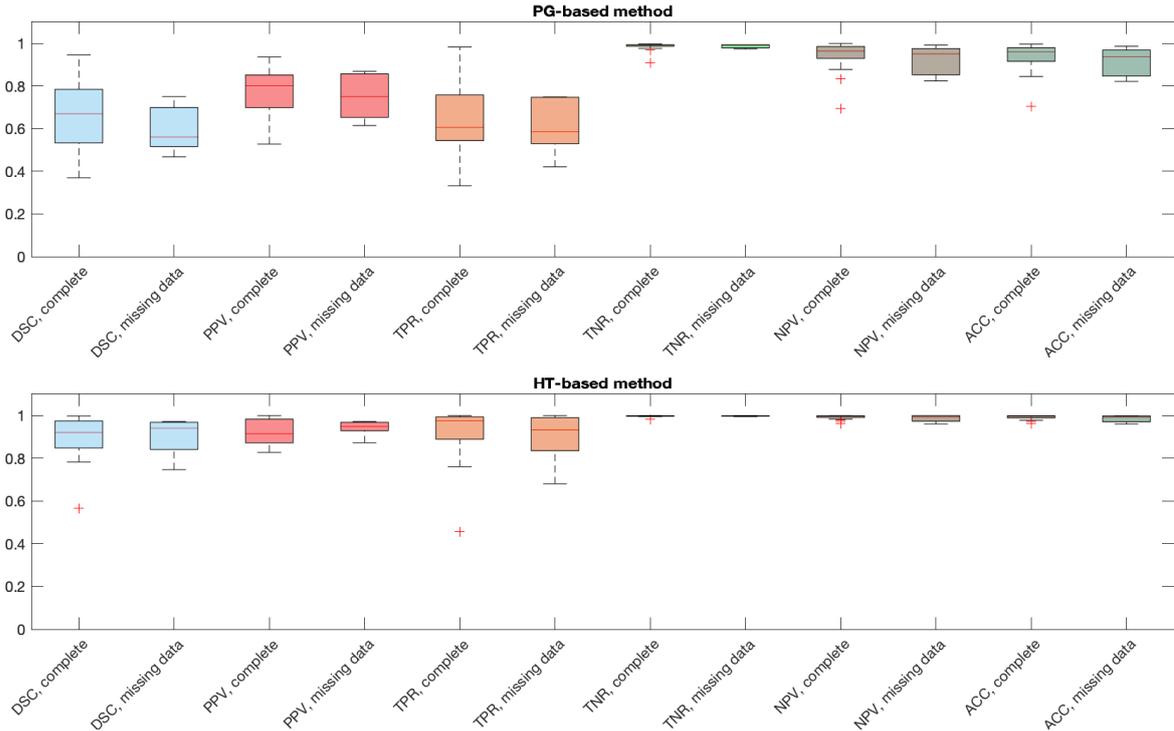
Figure 10: Performance of the PG- and HT-based methods, with an eye on models suffering from missing data. For these boxplots, we have made use of classification metrics presented in Table 2.

Computational time

All tests are performed on a desktop PC equipped with an Intel Core i9 processor (at 3.6 GHz) and a Windows 10 operating system. The routines have also been tested on a MacBook Pro equipped with macOS Catalina (version 10.15.7). We provide here some statistics of the execution times, obtained on the desktop PC:

- The PG-method has minimum, mean and maximum execution time corresponding to 1.7, 286.0 and 19074.0 seconds, respectively.

- The HT-method has minimum, mean and maximum execution time corresponding to 2.6, 50.7 and 358.0 seconds, respectively.

We observe that, for small point clouds, the PG-method is generally faster, while for big point clouds it is slower.

5 Conclusions

In this work we have proposed Fit4CAD, a benchmark for the evaluation and comparison of methods for fitting simple geometric primitives in point clouds representing CAD objects. The ground truth dataset of point clouds is segmented in geometric primitives and subdivided into a training set and a test set. In addition, a set of quality metrics and two fitting methods are given. In this work, evaluation metrics are used to quantify various performance aspects of geometric primitive fitting methods. In our intent, these metrics would assist both comparing with some methods in literature and allowing a parameters fine-tuning of a new method, in order to optimize it on a sufficiently large set of CAD models.

We hope the results of our comparison will inspire the development of new methods for primitive fitting, computational time being the main bottleneck in practice. In particular, it would be

Table 3: Approximation accuracy of the HT method.

	MFE	d_{dHaus}	d_1
PC 1	0.002	0.005	0.008
PC 2	0.007	0.011	0.149
PC 3	0.002	0.004	0.621
PC 4	0.002	0.003	0.000
PC 5	0.004	0.006	0.028
PC 6	0.003	0.005	0.001
PC 7	0.007	0.011	0.232
PC 8	0.004	0.013	0.001
PC 9	0.003	0.006	0.000
PC 10	0.005	0.019	1.203
PC 11	0.003	0.006	0.165
PC 12	0.002	0.004	0.058
PC 13	0.002	0.005	0.000
PC 14	0.004	0.006	1.264
PC 15	0.001	0.003	0.000
PC 16	0.003	0.004	0.029
PC 17	0.003	0.007	0.000
PC 18	0.003	0.004	0.324
PC 19	0.003	0.006	0.019
PC 20	0.002	0.003	0.001
PC 21	0.004	0.007	0.566
PC 22	0.002	0.005	0.000
PC 23	0.006	0.012	0.002
PC 24	0.001	0.003	0.000
PC 25	0.002	0.003	0.001
PC 26	0.001	0.003	0.000
PC 27	0.003	0.005	0.301
PC 28	0.001	0.002	0.000
PC 29	0.010	0.003	0.119
PC 30	0.002	0.009	0.003
PC 31	0.002	0.003	0.000
PC 32	0.006	0.007	0.000
PC 33	0.003	0.006	0.000
PC 34	0.003	0.005	0.004
PC 35	0.004	0.006	0.000

interesting to have a comparison with methods that use machine learning approaches, such as [30], because the dataset has been already organized in the form of a training set and a test set.

Regarding the two tested methods, the overall quality of the fitting is satisfactory for both. A rather unexpected conclusion is that over-segmentation is quite limited for both methods, while the combination of small and large primitives is a challenging task that often leads to a significant under-segmentation, see for instance the outcome on the model PC 9.

In future, we plan to continue to expand the dataset, even if we do not aim at a large scale dataset, for example by including more complex primitives and possibly considering specific contexts such as assembly models. Moreover, the flexibility of Fit4CAD permits the insertion of other available methods to reach a more complete view of the different typologies of approaches for geometric primitive fitting.

The benchmark is available at `https://github.com/chiararomanengo/Fit4CAD`.

Acknowledgments

References

[1] A. Kaiser, J. A. Ybanez Zepeda, and T. Boubekeur, "A Survey of Simple Geometric Primitives Detection Methods for Captured 3D Data," *Computer Graphics Forum*, vol. 38, no. 1, pp. 167–

196, 2019.

[2] A. Tal and E. Zuckerberger, "Mesh retrieval by components," in *Advances in Computer Graphics and Computer Vision* (J. e. a. Braz, ed.), pp. 44–57, Springer Berlin Heidelberg, 2007.

[3] C. R. Qi and et al., "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *2017 IEEE Conference CVPR*, pp. 77–85, 2017.

[4] V. Ganapathi-Subramanian and et al., "Parsing geometry using structure-aware shape templates," in *2018 Int. Conf. on 3D Vision (3DV)*, (Los Alamitos, CA, USA), pp. 672–681, IEEE Computer Society, sep 2018.

[5] A. F. Smeaton, P. Over, and W. Kraaij, "Evaluation campaigns and TRECVid," in *MIR '06: Proc. of the 8th ACM Int. Workshop on Multimedia Information Retrieval*, (New York, USA), pp. 321–330, ACM Press, 2006.

[6] C. S. Bojer and J. P. Meldgaard, "Kaggle forecasting competitions: An overlooked learning opportunity," *International Journal of Forecasting*, vol. 37, p. 587–603, Apr 2021.

[7] R. C. Veltkamp and et al., "SHREC2006: 3D Shape Retrieval Contest," Tech. Rep. UU-CS-2006-030, Dept. of Inf. and Comp. Sciences, Utrecht Univ., 2006.

[8] X. Chen, A. Golovinskiy, and T. Funkhouser, "A benchmark for 3D mesh segmentation," *ACM Trans. Graph.*, vol. 28, July 2009.

[9] G. Lavoué and et al., "SHREC'12 Track: 3D Mesh Segmentation," in *Eurographics Workshop on 3D Object Retrieval*, The Eurographics Association, 2012.

[10] Q. Zhou and A. Jacobson, "Thingi10k: A dataset of 10, 000 3d-printing models," *CoRR*, vol. abs/1605.04797, 2016.

[11] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.

[12] "The Shape Repository." http://visionair.ge.imati.cnr.it/ontologies/shapes/, 2011–2015.

[13] S. Koch, A. Matveev, Z. Jiang, F. Williams, A. Artemov, E. Burnaev, M. Alexa, D. Zorin, and D. Panozzo, "ABC: A big CAD model dataset for geometric deep learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[14] C. Geuzaine and J.-F. Remacle, "Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities," *Int. Journal for Num. Meth. in Eng.*, vol. 79, no. 11, pp. 1309–1331, 2009.

[15] A. Mathur, M. Pirron, and D. Zufferey, "Interactive Programming for Parametric CAD," *Computer Graphics Forum*, vol. 39, no. 6, pp. 408–425, 2020.

[16] T. Dokken, *Aspects of intersection algorithms and approximation.* PhD thesis, University of Oslo, 1997.

[17] O. J. D. Barrowclough and T. Dokken, "Approximate implicitization using linear algebra," *Journal of Applied Mathematics*, 2012.

[18] M. Kuhn and K. Johnson, *Applied Predictive Modeling.* Springer New York, 2018.

[19] M. M. Deza and E. Deza, *Encyclopedia of Distances.* Springer Berlin Heidelberg, 2009.

[20] P. V. C. Hough, "Method and means for recognizing complex patterns," 1962. US Patent 3,069,654.

[21] Y. Qie, L. Qiao, and N. Anwer, "Enhanced invariance class partitioning using discrete curvatures and conformal geometry," *Computer-Aided Design*, vol. 133, p. 102985, 2021.

[22] M. Pauly, M. Gross, and L. P. Kobbelt, "Efficient simplification of point-sampled surfaces," in *IEEE Visualization, 2002*, pp. 163–170, 2002.

[23] "Geometrical product specifications (GPS)– General concepts- Part 1: Model for geometrical specification and Verification," standard, ISO, Geneva, CH, Mar. 2011.

[24] N. Gelfand and L. J. Guibas, "Shape segmentation using local slippage analysis," in *Proceedings of the 2004 Eurographics – ACM SIGGRAPH Symposium on Geometry Processing*, SGP '04, (New York, NY, USA), p. 214–223, ACM, 2004.

[25] R. Schnabel, R. Wahl, and R. Klein, "Efficient RANSAC for Point-Cloud Shape Detection," *Computer Graphics Forum*, vol. 26, pp. 214–226, June 2007.

[26] M. C. Beltrametti and L. Robbiano, "An algebraic approach to Hough transforms," *Journal of Algebra*, vol. 37, pp. 669–681, 2012.

[27] M. Beltrametti, J. Sendra, J. Sendra, and M. Torrente, "Moore–Penrose approach in the Hough transform framework," *Applied Mathematics and Computation*, vol. 375, p. 125083, 2020.

[28] H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," *Discrete Comput. Geom.*, vol. 28, p. 511–533, Nov. 2002.

[29] S. Biasotti, A. Cerri, S. Pittaluga, D. Sobrero, and M. Spagnuolo, "Tracking the evolution of rainfall precipitation fields using persistent maxima," in *Proceedings of the Conference on Smart Tools and Applications in Computer Graphics*, pp. 29–37, 2016.

[30] S. Tulsiani, H. Su, L. J. Guibas, A. A. Efros, and J. Malik, "Learning shape abstractions by assembling volumetric primitives," in *Computer Vision and Pattern Regognition (CVPR)*, 2017.

A   File type description

The geometrical information of the primitives are provided in the files "*PCi_parametric*" and "*PCi_implicit*". Here, we describe in detail the equations of the parametric and implicit representation that they contain. Notice that, although the simple primitives shapes considered in this benchmark are polynomial (i.e., can be written as the zero set of a bivariate polynomial), the parametric representation we provide for cylinders, cones, spheres and tori are written in terms of trigonometric functions.

A.1   Parametric representations

- Plane:

$$\begin{cases} x = a_1 u + b_1 v + c_1 \\ y = a_2 u + b_2 v + c_2 \\ z = a_3 u + b_3 v + c_3 \end{cases}$$

  The parameters for a plane are stored as follows:

  ```
  [Plane, [a1 a2 a3 b1 b2 b3 c1 c2 c3]]
  ```

- Cylinder:

$$\begin{cases} x = a_1 \cos(u) + b_1 \sin(u) + c_1 v + d_1 \\ y = a_2 \cos(u) + b_2 \sin(u) + c_2 v + d_2 \\ z = a_3 \cos(u) + b_3 \sin(u) + c_3 v + d_3 \end{cases}$$

  The parameters for a cylinder are stored as follows:

```
[Cylinder, [a1 a2 a3 b1 b2 b3 c1 c2 c3 d1 d2 d3]
```

- Cone:

$$\begin{cases} x = a_1 \cos(u) + b_1 \sin(u) + c_1 v \cos(u) + d_1 v \sin(u) + e_1 v + f_1 \\ y = a_2 \cos(u) + b_2 \sin(u) + c_2 v \cos(u) + d_2 v \sin(u) + e_2 v + f_2 \\ z = a_3 \cos(u) + b_3 \sin(u) + c_3 v \cos(u) + d_3 v \sin(u) + e_3 v + f_3 \end{cases}$$

The parameters for a cone are stored as follows:

```
[Cone, [a1 a2 a3 b1 b2 b3 c1 c2 c3 d1 d2 d3 e1 e2 e3 f1 f2 f3]]
```

- Sphere:

$$\begin{cases} x = a_1 \cos(u) \cos(v) + b_1 \sin(u) \cos(v) + c_1 \sin(v) + d_1 \\ y = a_2 \cos(u) \cos(v) + b_2 \sin(u) \cos(v) + c_2 \sin(v) + d_2 \\ z = a_3 \cos(u) \cos(v) + b_3 \sin(u) \cos(v) + c_3 \sin(v) + d_3 \end{cases}$$

The parameters for a sphere are stored as follows:

```
[Sphere, [a1 a2 a3 b1 b2 b3 c1 c2 c3 d1 d2 d3]]
```

- Torus:

$$\begin{cases} x = a_1 \cos(u) + b_1 \sin(u) + c_1 \cos(u) \cos(v) + d_1 \sin(u) \cos(v) + e_1 \sin(v) + f_1 \\ y = a_2 \cos(u) + b_2 \sin(u) + c_2 \cos(u) \cos(v) + d_2 \sin(u) \cos(v) + e_2 \sin(v) + f_2 \\ z = a_3 \cos(u) + b_3 \sin(u) + c_3 \cos(u) \cos(v) + d_3 \sin(u) \cos(v) + e_3 \sin(v) + f_3 \end{cases}$$

The parameters for a torus are stored as follows:

```
[Torus, [a1 a2 a3 b1 b2 b3 c1 c2 c3 d1 d2 d3 e1 e2 e3 f1 f2 f3]]
```

A.2 Implicit representations

- Plane:
$$ax + by + cz + d = 0$$
The coefficients for a plane are stored as follows:

```
[Plane, [a b c d]]
```

- Cylinder:
$$ax^2 + by^2 + cy^2 + 2(dxy + exz + fyz) + 2(gx + hy + iz) + l = 0$$
The coefficients for a cylinder are stored as follows:

```
[Cylinder,[a b c d e f g h i l]]
```

- Cone:
$$ax^2 + by^2 + cy^2 + 2(dxy + exz + fyz) + 2(gx + hy + iz) + l = 0$$
The coefficients for a cone are stored as follows:

```
[Cone, [a b c d e f g h i l]]
```

- Sphere:
$$ax^2 + by^2 + cy^2 + 2(dxy + exz + fyz) + 2(gx + hy + iz) + l = 0$$
The coefficients for a sphere are stored as follows:

```
[Sphere, [a b c d e f g h i l]]
```

- Torus: the coefficients of the implicit representation are provided in the form of a polynomial of degree 4 in x, y and z; they are stored in reverse lexicographic order.

B   Results of the test set segmentations
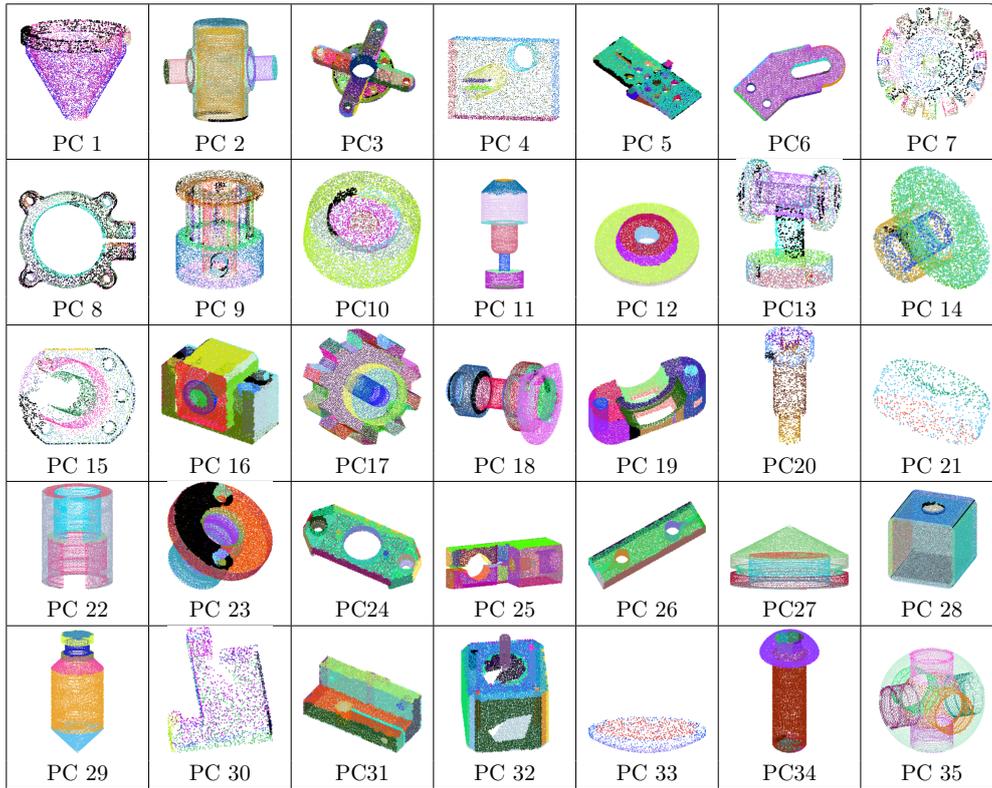
Figure 11: Segmentations obtained via the PG-based method.
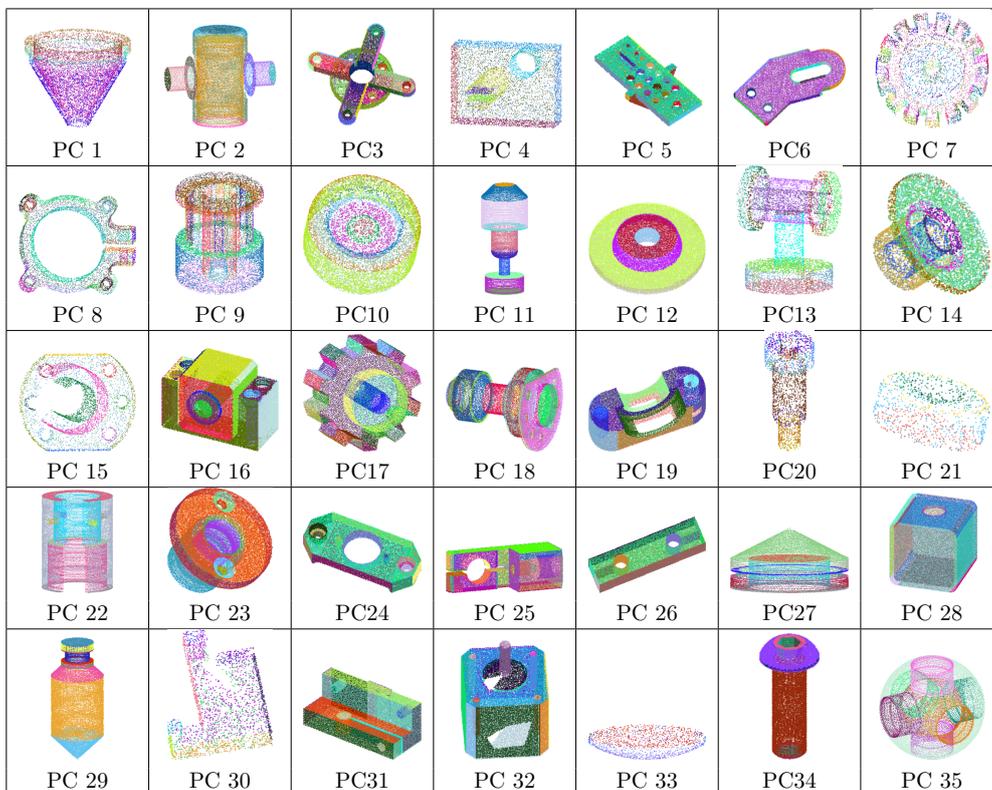
Figure 12: Segmentations obtained via the Hough-based method.