Supplementary Material

CADOps-Net: Jointly Learning CAD Operation Types and Steps from Boundary-Representations

A. CC3D-Ops Dataset

In this section, we provide more details on the proposed *CC3D-Ops* dataset. First, the method used to extract the *op.step* and *op.type* will be briefly discussed. Then, statistics demonstrating the complexity of the models and showing the distribution of the labels will be presented.

A.1. CC3D-Ops Label Extraction

The proposed *CC3D-Ops* dataset contains B-Reps with per-face *op.type* and *op.step* annotations. An important aspect of segmenting faces into different construction steps of modeling operations is that these labels come from the real construction history of each CAD model in the dataset. In our case, this information is obtained from the native SolidWorks [2] Part File (.sldprt) format of a CAD model. A set of tools were developed based on the Solidworks API [2] to traverse a CAD model's construction history and to assign each face generated by respective modeling operation its *op.type* and *op.step* labels in B-Rep.

A.2. Statistics

Model complexity: From the sample *CC3D-Ops* CAD models (in B-Rep format) displayed in Figure 1, it can be noted that *CC3D-Ops* offers a wide variety of models both in terms of complexity and category. Figure 2 shows the distribution of the number of faces per model for the *CC3D-Ops* and Fusion360 [3] datasets as box plots. This figure shows that the models in *CC3D-Ops* generally have more faces than in Fusion360. While 90% of the models of the Fusion360 dataset have 30 faces or less, such models represent only 50% of *CC3D-Ops*. This difference between the two datasets is further demonstrated in Figure 3, where it can be observed that the models in *CC3D-Ops* tend to be made of more CAD operation steps than for Fusion360.

CAD Operation Type Labels: The *op.type* face labels indicate the type of CAD operation used during the design process. While the most common CAD operation types (such as *extrusion*, *fillet* ...) are shared among most CAD software applications, some are software specific. The *CC3D-Ops* dataset contains 11 different *op.type* labels: *extrude side*, *extrude end*, *revolve side*, *revolve end*, *cut extrude side*, *cut extrude end*, *cut revolve side*, *cut revolve end*, *fillet*, *chamfer* and *other*. The *other op.type* represents less common types such as *helix*, *sweep*, *dome*, *etc*. The bar chart in Figure 4 displays the number of faces for each *op.type* label. The two least common *op.type* labels are *revolve end* and *cut revolve end* and the two most common operation types are *extrude side* and *other*. For a comparison with the Fusion360 dataset, we refer the reader to [1] where a similar bar chart can be found.

B. Further Experimental Analysis

In this section, we first analyze quantitative results on the per class IoUs for the *op.type* segmentation task. Then, some visual examples of the *CADOps-Net* predictions are presented.

B.1. CAD Operation Type IoUs

As in [1], an analysis of each *op.type* IoU is presented. Table 1 and 2 show the IoU results of each *op.type* label for the Fusion360 and *CADOps-Net* datasets respectively. The results are shown for both *CADOps-Net* without joint learning (*Ours w/o* JL^{-}) and with joint learning (*Ours w/* JL^{+}). As noted in Section 6.2 of the main paper, the joint learning strategy does not have a significant impact on the *op.type* predictions. This is particularly the case on the Fusion360 dataset as shown in Table 1. *Ours w/* JL^{+} achieves slightly higher IoU for each class. On the other hand, the same trend cannot be found in the results obtained from the *CC3D-Ops* dataset. For 6 out of the 11 *op.type* classes, the difference between the IoUs obtained with joint learning and without is relatively small (less than 2%). For the *op.type* classes *cut revolve side* and *chamfer*, *Ours w/o* JL^{-} scores higher than the *Ours w/* JL^{+} by 3.9% and 5.5% respectively. However, the joint learning method achieves higher results on 3 out of the 4 least common classes, namely *cut revolve end*, *revolve end* and *cut extrude end*. In particular, for the *revolve end op.type* that represents 0.17% of the dataset, the joint learning strategy results in an IoU that is about 17% higher than without joint learning. This demonstrates that even if the joint learning strategy achieves a comparable mIoU as without joint learning, *Ours w/* JL^{+} is able to learn more meaningful features for the underrepresented *op.types*.

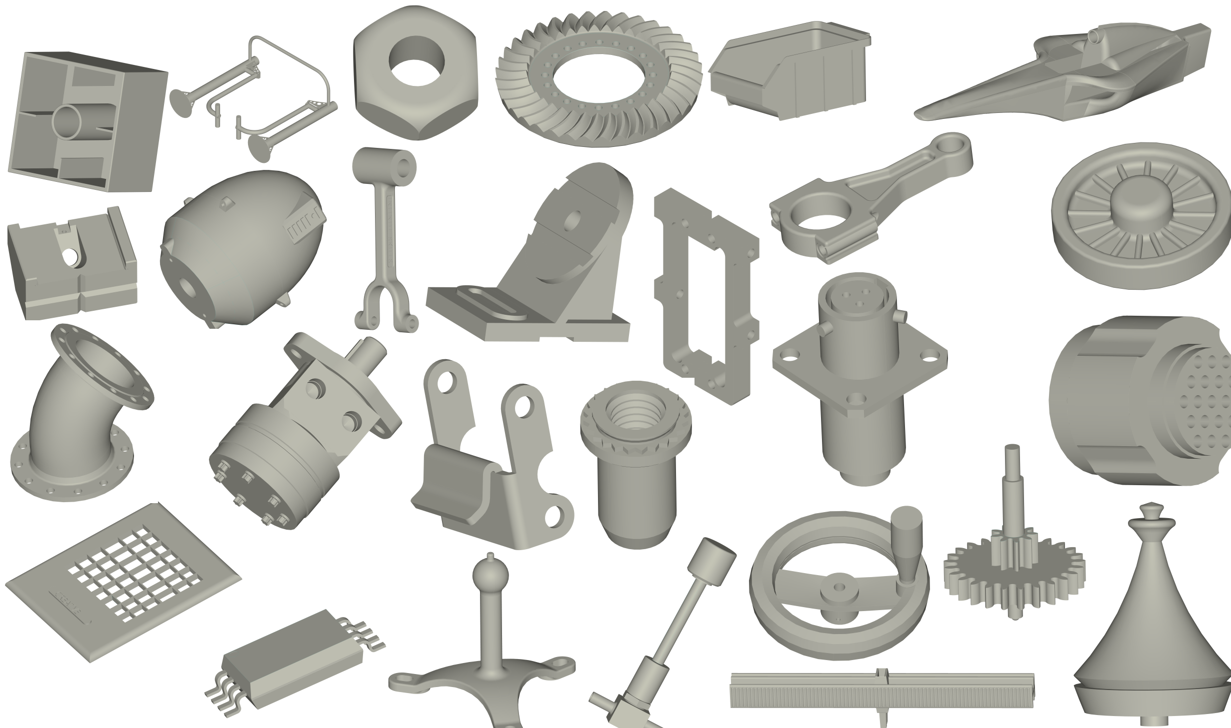

Figure 1. Sample CAD models from the *CC3D-Ops* dataset.

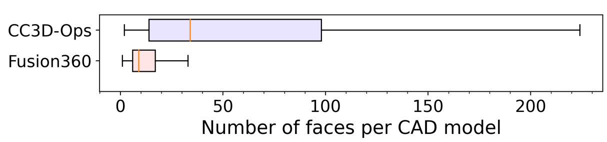

Figure 2. Box plot showing the distribution of models for the *CC3D-Ops* and Fusion360 [3] datasets with respect to the number of faces per model.

Fusion360	Per class IoU	
	Ours w/o JL$^-$	*Ours w/ JL*$^+$
Extrude side	94.0	**94.6**
Extrude end	91.7	**92.4**
Cut side	82.1	**83.9**
Cut end	75.2	**77.1**
Revolve side	85.1	**86.5**
Revolve end	48.7	**48.9**
Chamfer	91.2	**92.1**
Fillet	97.6	**97.8**

Table 1. *op.type* per class IoU for the Fusion360 dataset. All results are expressed as percentages.

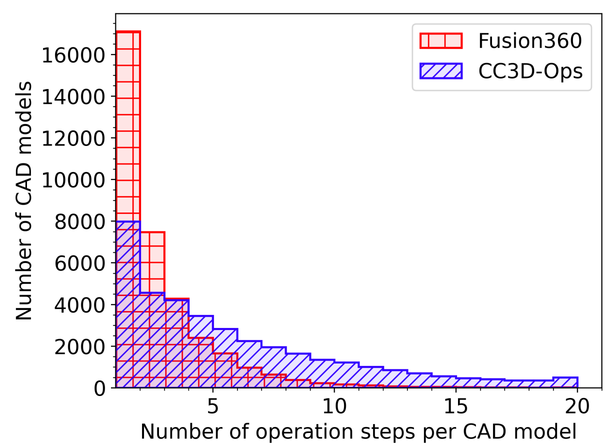

Figure 3. Histogram showing the distribution of models for the *CC3D-Ops* and Fusion360 [3] datasets with respect to the number of *op.steps* per model. Models with a number of *op.steps* between 0 and 20 represent over 96% of *CC3D-Ops*.

B.2. CAD Operation Types Qualitative Results

Figure 5 shows sample results from *CADOps-Net* for the *op.type* segmentation task on both the Fusion360 and *CC3D-Ops* datasets. As discussed in Section 6.2 of the main paper, the *op.type* accuracy is not correlated to the complexity of the models. In particular, this can be observed from the results presented on the Fusion360 dataset in which *CADOps-Net* sometimes fails to predict the correct *op.types* for some simple models. Despite not matching the ground truth, some predictions are still valid within the context of CAD modelling, as outlined in Section 6.5.

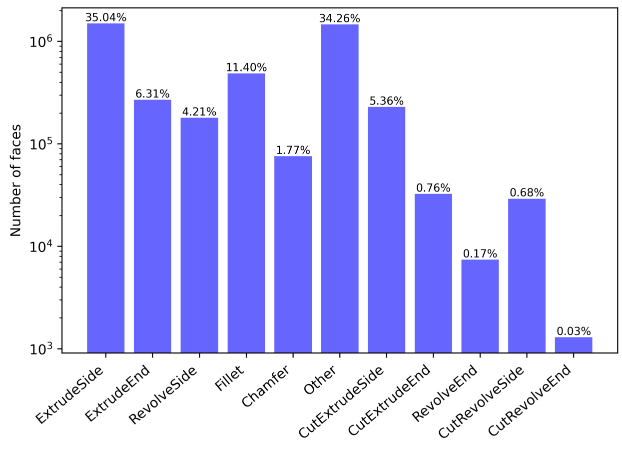

Figure 4. Bar graph of the number of faces for each *op.type* label over the *CC3D-Ops* dataset. The numbers above each bar represents the percentage of the number of faces with the corresponding type. Note: a *log* scale is used for the vertical axis.

CC3D-Ops	Per class IoU	
	Ours w/o JL$^-$	*Ours w/ JL*$^+$
Extrude side	**65.4**	64.9
Extrude end	59.7	**60.2**
Cut extrude side	17.8	**18.1**
Cut extrude end	10.0	**15.4**
Cut revolve side	**22.2**	18.3
Cut revolve end	1.1	**4.6**
Revolve side	**60.3**	59.8
Revolve end	23.8	**41.2**
Chamfer	**69.6**	64.4
Fillet	**84.1**	83.1
Other	**58.5**	57.2

Table 2. *op.type* per class IoU for the *CC3D-Ops* dataset. All results are expressed as percentages.

B.3. CAD Operation Step Qualitative Results

Figure 6 displays qualitative results for the *op.step* segmentation task on the *CC3D-Ops* and Fusion360 datasets. These qualitative results illustrate that while *CADOps-Net* is able to make accurate predictions for models with a small number of *op.steps*, the *op.step* accuracy decreases as the number of *op.steps* increases, as explained in Section 6.2.

C. Sketch Recovery Process and Examples

In CAD modeling, sketches are considered as starting points to build a kernel structure (in our case B-Rep) of the desired solid model. Nevertheless, the sketches are only part of the forward design process, and tracing them back from the final B-Rep is not straightforward. In this section, we provide more details about the process of sketch recovery introduced in Section 6.4 and present more qualitative results.

Algorithm 1: Sketch Recovery Algorithm

Input: $\mathcal{B}, \widehat{\mathbf{S}} \in [0,1]^{N_f \times k_s}, \widehat{\mathbf{T}} \in [0,1]^{N_f \times k_t}$
Output: a set of sketches $\mathbb{Q}$

1 **foreach** *id* $\mathbf{s} \in \{\text{argmax}\, \widehat{\mathbf{S}}_{j:},\ \forall j \in [1,\dots,N_f]\}$ **do**
2   **foreach** *id* $\mathbf{t} \in \{\text{argmax}\, \widehat{\mathbf{T}}_{j:},\ \forall j \in [1,\dots,N_f]\}$ **do**
3     $\mathcal{F} \leftarrow \{.., f_{\mathbf{s}}, ...\}$; `/* Group face ids` $(\mathbf{s},\mathbf{t})$ `into` $f_{\mathbf{s}}$ `of` *`op.type`* `'extrude side' at different` *`op.steps`* `*/`
4   **end**
5 **end**
6 **foreach** $f_{\mathbf{s}} \in \mathcal{F}$ **do**
7   compute normals $\{\hat{\mathbf{n}}_i^{\mathbf{s}}, \forall i \in f_{\mathbf{s}}\}$;
8   $\hat{o}_{\mathbf{s}} \leftarrow$ centroid of merged faces $\in f_{\mathbf{s}}$;
9   $\hat{a}_{\mathbf{s}} \leftarrow$ optimal extrusion axis $\sum_{\substack{i,j \in f_{\mathbf{s}} \\ i \neq j}} \frac{\hat{\mathbf{n}}_i^{\mathbf{s}} \times \hat{\mathbf{n}}_j^{\mathbf{s}}}{|\hat{\mathbf{n}}_i^{\mathbf{s}} \times \hat{\mathbf{n}}_j^{\mathbf{s}}|}$;
10   $\widehat{\mathbb{Q}}^{\mathbf{s}} \leftarrow$ project merged faces $\in f_{\mathbf{s}}$ on a plane with center $\hat{o}_{\mathbf{s}}$ and axis $\hat{a}_{\mathbf{s}}$
11 **end**

Proposition 1 *If we assume that not all faces of a B-Rep are either orthogonal or parallel to each other, then the base sketch-profile of the merged faces may not be co-planar, and therefore extrusion or revolution axis may not be orthogonal to the normals of co-faces.*

Algorithm 1 describes the process for retrieving sketches at different *op.steps* using *CADOps-Net* predictions ($\widehat{\mathbf{S}}$ and $\widehat{\mathbf{T}}$), assuming the proposition 1 holds over a B-Rep when extrusion-only operation is involved. In particular, *CADOps-Net* predictions of *op.steps* are used to group the faces of the B-Rep that were created by a single sketch. Among these faces, the ones created by *extrude side* are identified by the *op.type* predictions (lines 1 to 5 of Algorithm 1). These faces are denoted $f_{\mathbf{s}}$ and are considered to compute the extrusion axis and the centroid, then the projection plane as described in lines 6 to 11 of Algorithm 1.

Figure 7 and 8 illustrate the qualitative results of sketches recovered using *CADOps-Net* predictions and Algorithm 1 on randomly selected samples made by extrusions from Fusion360 [3] dataset. We dissect the sketch results based on correctly and incorrectly predicted *op.steps* in the two figures. We can observe that the sketch recovery is successful when the predictions of *op.steps* are correct (Figure 7). In Figure 8, the B-Reps were segmented into two *op.steps*, while the ground truth annotations indicated that they were designed through three *op.steps*. Such incorrect *op.step* prediction impacted the sketch recovery and resulted in erroneous sketches.

References

[1] Joseph G. Lambourne, Karl D.D. Willis, Pradeep Kumar Jayaraman, Aditya Sanghi, Peter Meltzer, and Hooman Shayani. Brepnet: A topological message passing system for solid models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12773–12782, June 2021. 1

[2] Solidworks. 3D CAD Design Software. https://www.solidworks.com/. Online: accessed 02-June-2022. 1

[3] Karl D. D. Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G. Lambourne, Armando Solar-Lezama, and Wojciech Matusik. Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. *ACM Trans. Graph.*, 40(4), jul 2021. 1, 2, 3

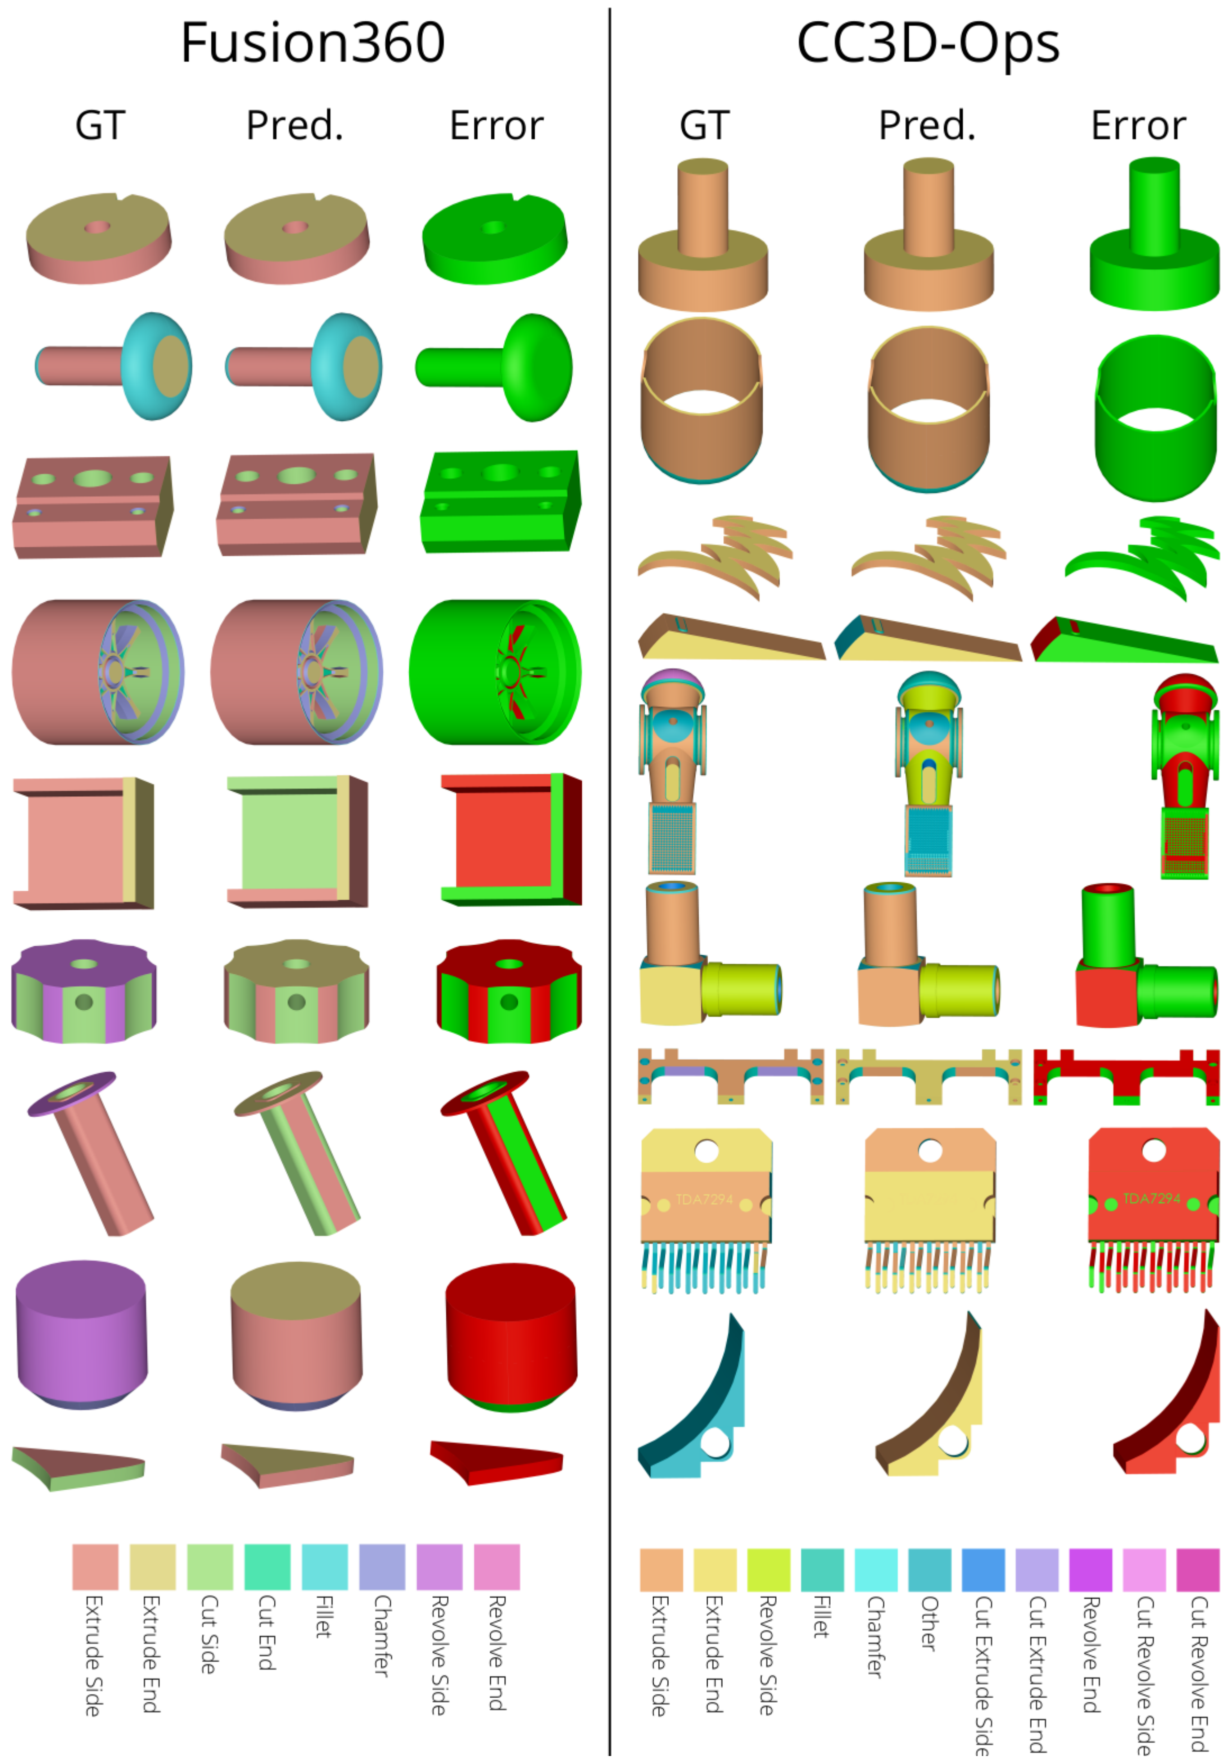

Figure 5. CAD operation step qualitative results on the Fusion360 (left) and *CC3D-Ops* (right) datasets. For the *CADOps-Net* ground truth (GT) and predictions (Pred.). Correctly segmented faces are shown in green and incorrect in red in the Error columns.

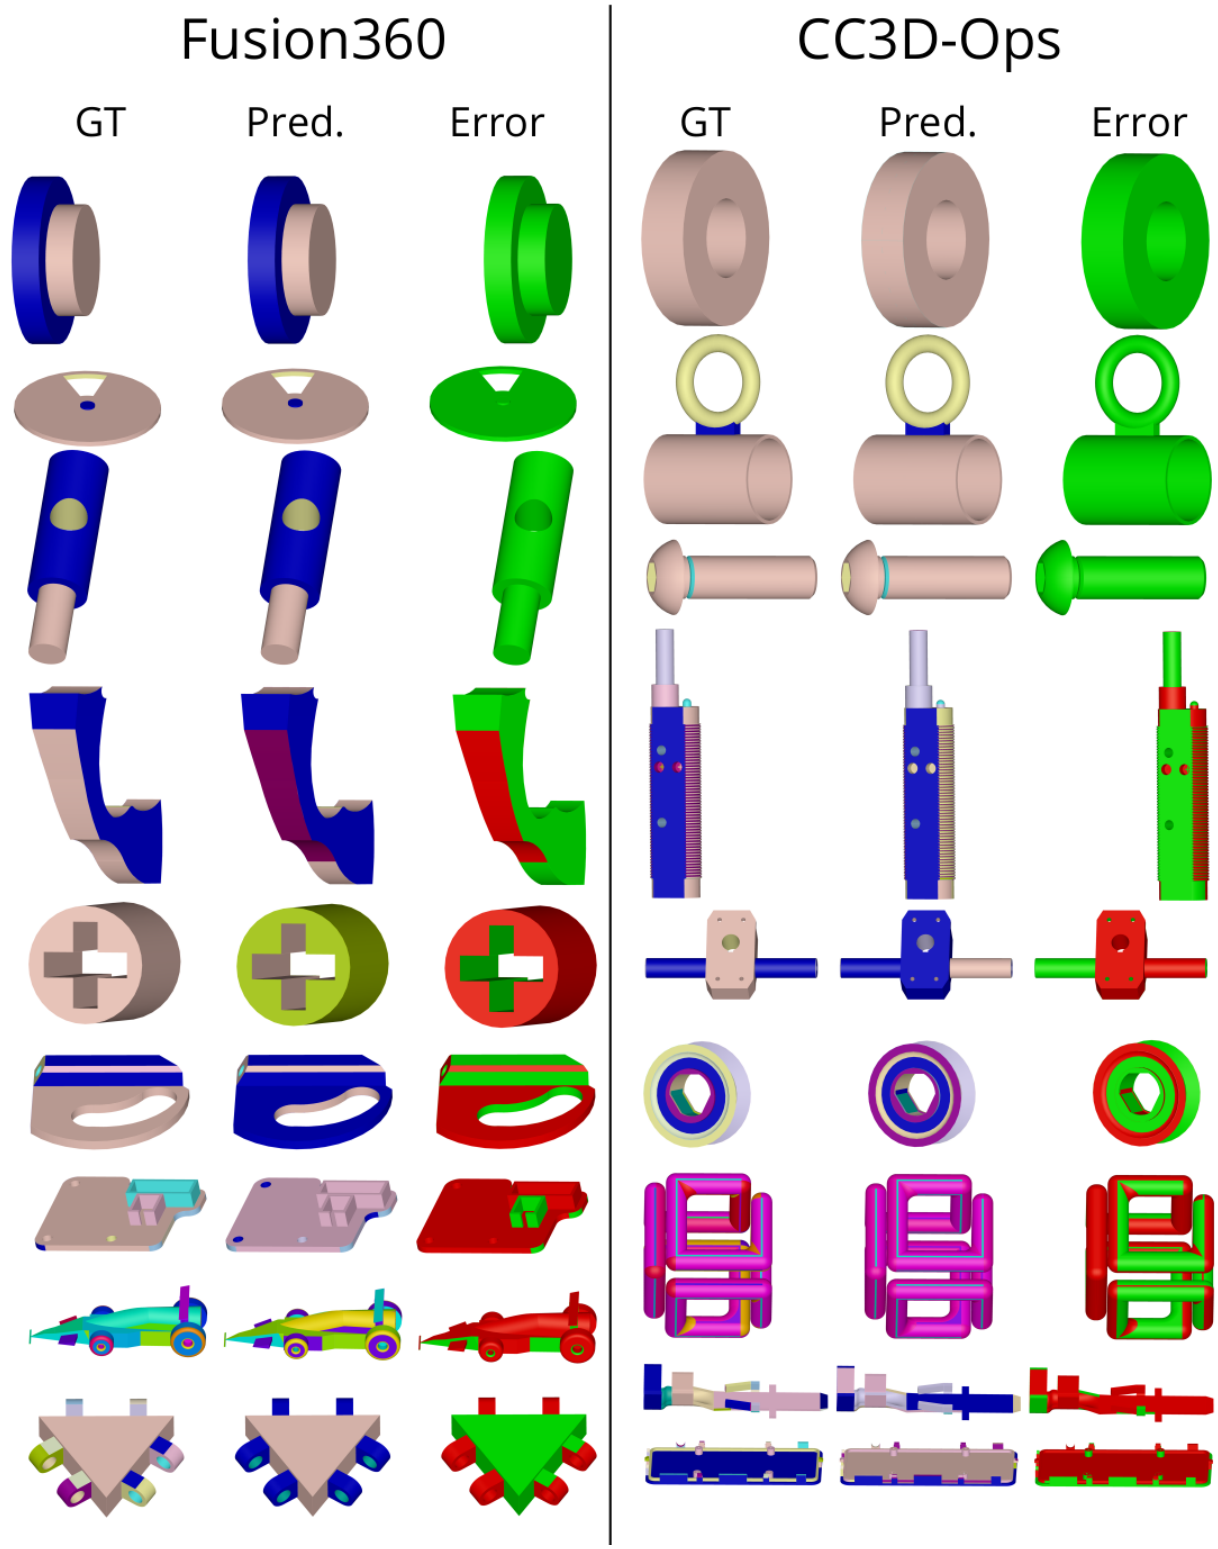

Figure 6. CAD operation step qualitative results on the Fusion360 (left) and *CC3D-Ops* (right) datasets.For the *CADOps-Net* ground truth (GT) and predictions (Pred.), each color represents a CAD operation step. Correctly segmented faces are shown in green and incorrect in red in the Error columns.

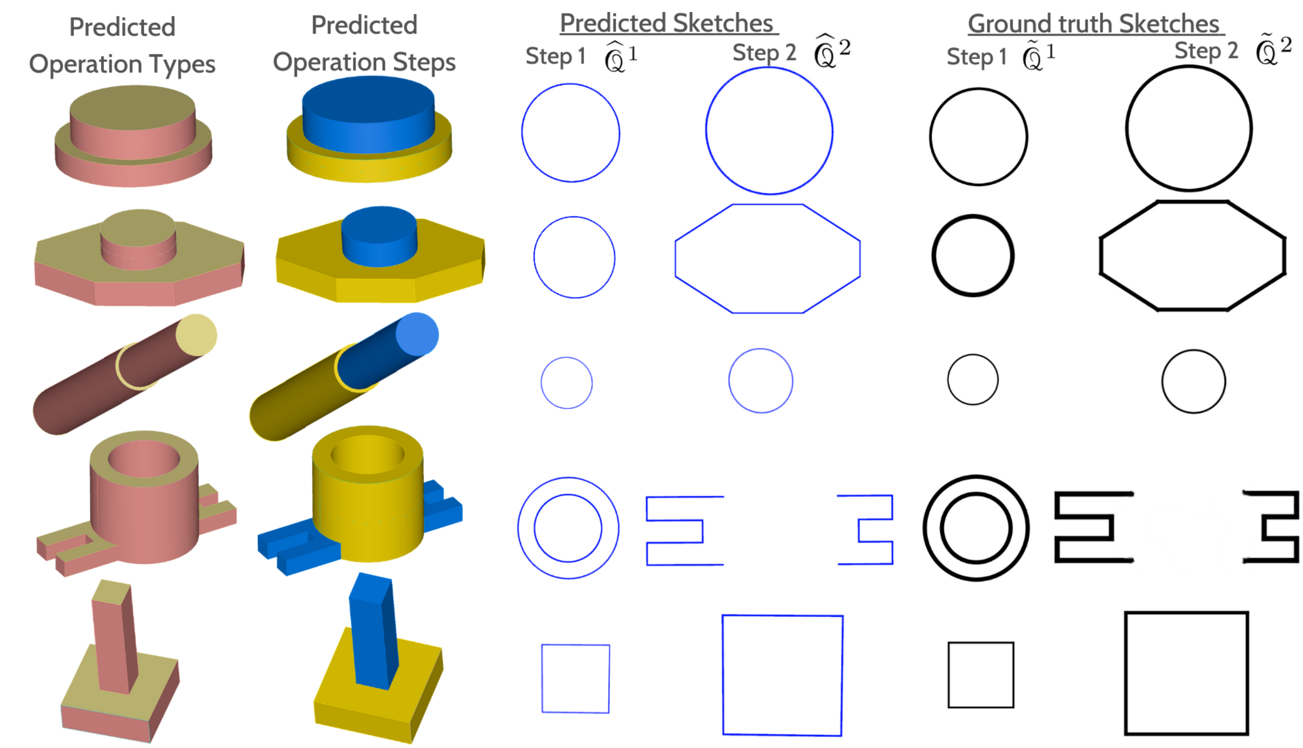

Figure 7. Qualitative results on sketch recovery from correctly predicted *op.types* and *op.steps* by CADOps-Net. The models shown above include exactly two operation steps.

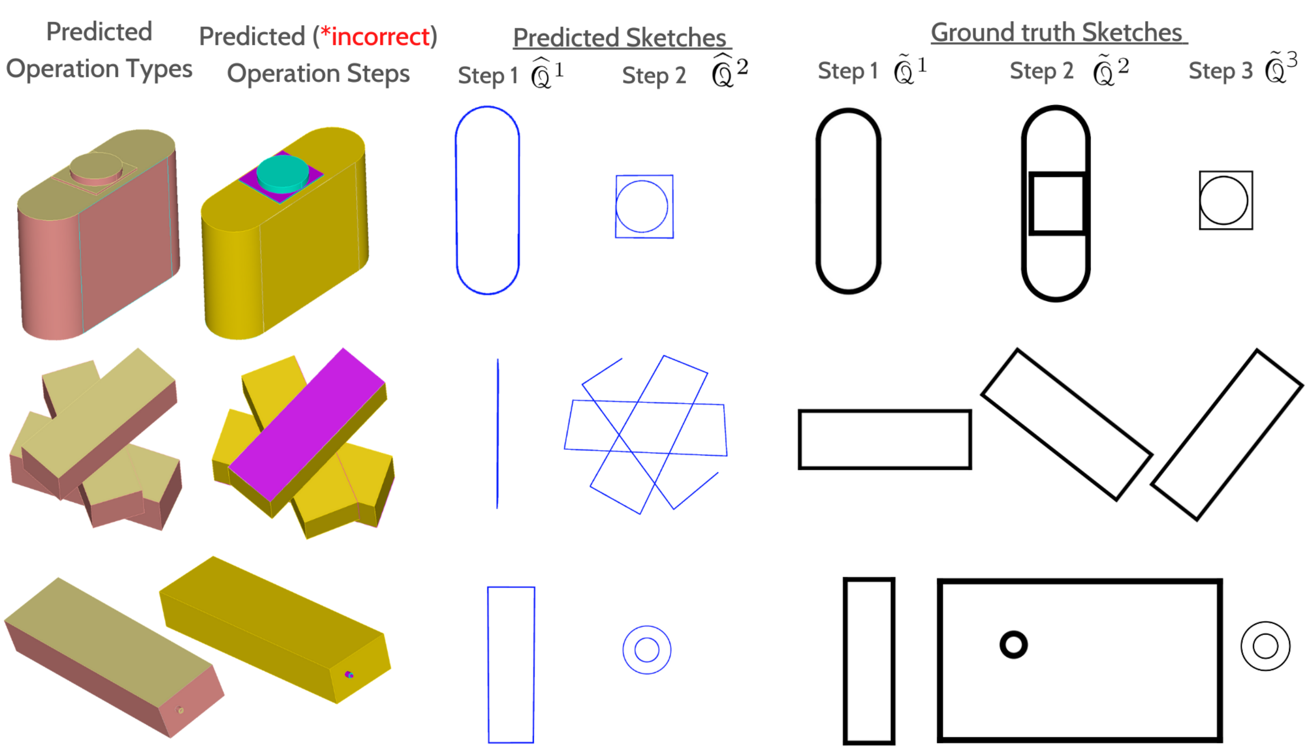

Figure 8. Qualitative results on sketch recovery from incorrectly predicted *op.types* and *op.steps* by CADOps-Net. The models shown above include exactly three operation steps.