Table 1: *Description of the online photo dataset used. The number of labeled photos and the number of semantic labels in each category are listed. The number of models tested and their locations in the paper and the supplementary material are also provided.*

Category	# semantic labels	# labeled photos	# tested models	In paper	In supp. material
Chair	4	500	10	Figs. 2 & 13	Fig. **??** & **??**
Baby Stroller	6	400	3	Fig. 1	Fig. **??**
Truck	3	400	5	Fig. 12	Fig. **??**
Lamp	3	344	6	Fig. 12	Fig. **??**
Vase	4	300	5		Fig. **??**
Table	2	250	3		Fig. **??**
Bike	5	181	6	Fig. 12 & 14	Fig. **??**
Pavilion	3	60	4	Fig. 12	Fig. **??**
Guitar	3	20	3		Fig. **??**
Fourleg	5	234	4		Fig. **??**
Robot	4	174	3	Fig. 12	Fig. **??**

Table 2: *Confusion matrix between different shape categories generated by different approaches. The testing dataset consists 400 shapes, 50 for each category (randomly sampled from the first eight categories in Table* **??***). Each shape is matched to all others in the database and the top five matches are used to compute the labeling probability. Direct comparison among different approaches is given in the next table.*

(a) Our Approach.

	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp
Stroller	**0.74**	0.01	0.13	0.02	0.02	0.04	0.03	0.01
Bike	0.00	**0.99**	0.00	0.00	0.01	0.00	0.00	0.00
Chair	0.09	0.00	**0.74**	0.03	0.03	0.03	0.04	0.04
Pavilion	0.01	0.02	0.02	**0.84**	0.08	0.02	0.01	0.00
Table	0.00	0.00	0.01	0.01	**0.93**	0.04	0.01	0.00
Truck	0.00	0.00	0.00	0.20	0.05	**0.90**	0.03	0.00
Vase	0.01	0.00	0.02	0.04	0.00	0.01	**0.86**	0.06
Lamp	0.02	0.01	0.04	0.00	0.02	0.02	0.08	**0.81**

(b) Inner-distance shape context (IDSC).

	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp
Stroller	**0.87**	0.02	0.00	0.00	0.02	0.00	0.02	0.00
Bike	0.00	**0.95**	0.01	0.00	0.02	0.00	0.00	0.02
Chair	0.01	0.05	**0.67**	0.04	0.13	0.00	0.03	0.07
Pavilion	0.00	0.00	0.04	**0.87**	0.06	0.00	0.02	0.01
Table	0.00	0.01	0.02	0.03	**0.88**	0.03	0.01	0.02
Truck	0.00	0.00	0.00	0.00	0.01	**0.91**	0.07	0.01
Vase	0.06	0.00	0.01	0.00	0.02	0.07	**0.80**	0.04
Lamp	0.01	0.01	0.02	0.00	0.05	0.00	0.01	**0.90**

(c) GIST descriptor

	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp
Stroller	**0.87**	0.00	0.06	0.01	0.01	0.02	0.02	0.01
Bike	0.04	**0.96**	0.00	0.00	0.00	0.00	0.00	0.00
Chair	0.10	0.00	**0.64**	0.07	0.02	0.02	0.07	0.08
Pavilion	0.00	0.00	0.03	**0.82**	0.08	0.06	0.00	0.01
Table	0.01	0.00	0.00	0.01	**0.89**	0.08	0.01	0.00
Truck	0.00	0.00	0.00	0.00	0.01	**0.98**	0.02	0.00
Vase	0.02	0.00	0.01	0.00	0.01	0.07	**0.78**	0.11
Lamp	0.02	0.00	0.03	0.00	0.03	0.07	0.12	**0.73**

(d) Light field descriptor (LFD)

	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp
Stroller	**0.51**	0.06	0.05	0.10	0.09	0.05	0.01	0.13
Bike	0.11	**0.74**	0.04	0.01	0.03	0.00	0.01	0.06
Chair	0.09	0.09	**0.40**	0.04	0.04	0.10	0.11	0.13
Pavilion	0.08	0.01	0.03	**0.45**	0.11	0.24	0.06	0.02
Table	0.08	0.01	0.04	0.07	**0.58**	0.10	0.04	0.08
Truck	0.00	0.00	0.02	0.17	0.04	**0.60**	0.09	0.08
Vase	0.02	0.00	0.08	0.12	0.05	**0.31**	0.30	0.12
Lamp	0.07	0.05	0.09	0.02	0.10	0.12	0.12	**0.43**

Table 3: *The retrieval rate for different shape categories generated by different approaches. The same testing dataset as in Table* **??** *is used. Depending on the setting, the top 1, 2, 5, and 10 matches for each shape in the dataset are used to compute the retrieval rates. The highest retrieval rate under each setting is shown in **bold**. The comparison suggests that our approach offers better average retrieval accuracy when only one or two shapes are returned, but is outperformed by existing approaches when more shapes are required. This is likely due to the piecewise linear scaling we used, which could scale shapes from different categories and make them appear similar.*

(a) Top 1 match.

Algorithm	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp	Average
Ours	41/50	**50/50**	**47/50**	46/50	**47/50**	48/50	**49/50**	43/50	**0.93**
IDSC	45/50	49/50	40/50	**49/50**	46/50	**49/50**	45/50	**46/50**	0.92
GIST	**48/50**	**50/50**	34/50	45/50	46/50	**49/50**	43/50	42/50	0.89
LFD	31/50	40/50	17/50	24/50	32/50	36/50	16/50	28/50	0.56

(b) Top 2 matches.

Algorithm	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp	Average
Ours	76/100	**99/100**	**88/100**	89/100	**95/100**	93/100	**94/100**	84/100	**0.90**
IDSC	88/100	95/100	79/100	**96/100**	88/100	98/100	84/100	**91/100**	**0.90**
GIST	**92/100**	97/100	68/100	90/100	91/100	**99/100**	86/100	80/100	0.88
LFD	54/100	77/100	41/100	47/100	63/100	67/100	34/100	46/100	0.54

(c) Top 5 matches.

Algorithm	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp	Average
Ours	185/250	**247/250**	**182/250**	209/250	**232/250**	226/250	**214/250**	202/250	0.85
IDSC	217/250	237/250	168/250	**217/250**	221/250	228/250	199/250	**226/250**	**0.86**
GIST	**218/250**	239/250	160/250	207/250	222/250	**245/250**	196/250	182/250	0.83
LFD	128/250	185/250	99/250	112/250	144/250	148/250	73/250	108/250	0.50

(d) Top 10 matches.

Algorithm	Stroller	Bike	Chair	Pavilion	Table	Truck	Vase	Lamp	Average
Ours	291/500	**492/500**	**333/500**	**410/500**	393/500	375/500	**385/500**	330/500	0.75
IDSC	**416/500**	474/500	305/500	401/500	**420/500**	422/500	345/500	**430/500**	0.79
GIST	399/500	464/500	280/500	401/500	**420/500**	**471/500**	339/500	340/500	**0.80**
LFD	229/500	354/500	166/500	198/500	265/500	264/500	154/500	194/500	0.46

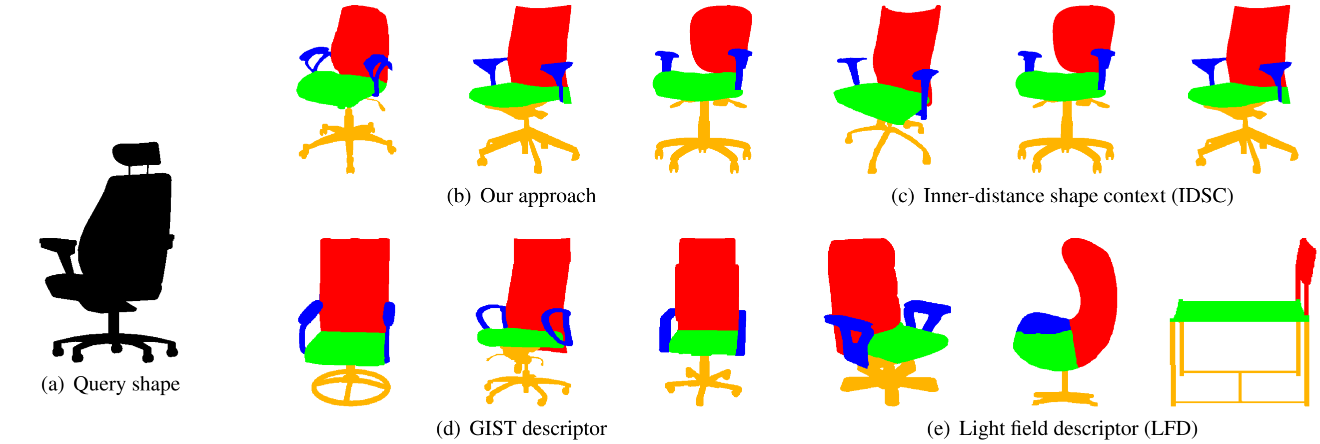

Figure 1: *Visual comparison for shapes retrieved by different approaches for a chair projection. This is the same figure as Figure 9 in the paper, but with bigger size for better visibility.*

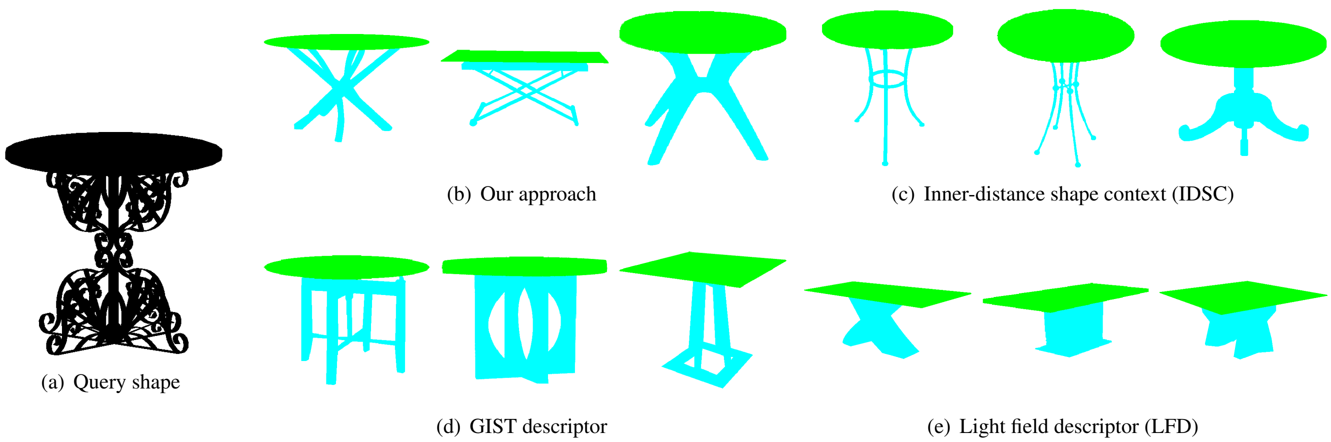

Figure 2: *Visual comparison for shapes retrieved by different approaches for a table projection (a). Our approach (b) returns shapes that have similar overall topology but much simpler in details, suggesting the BiSH distance measure is robust against small internal holes.*

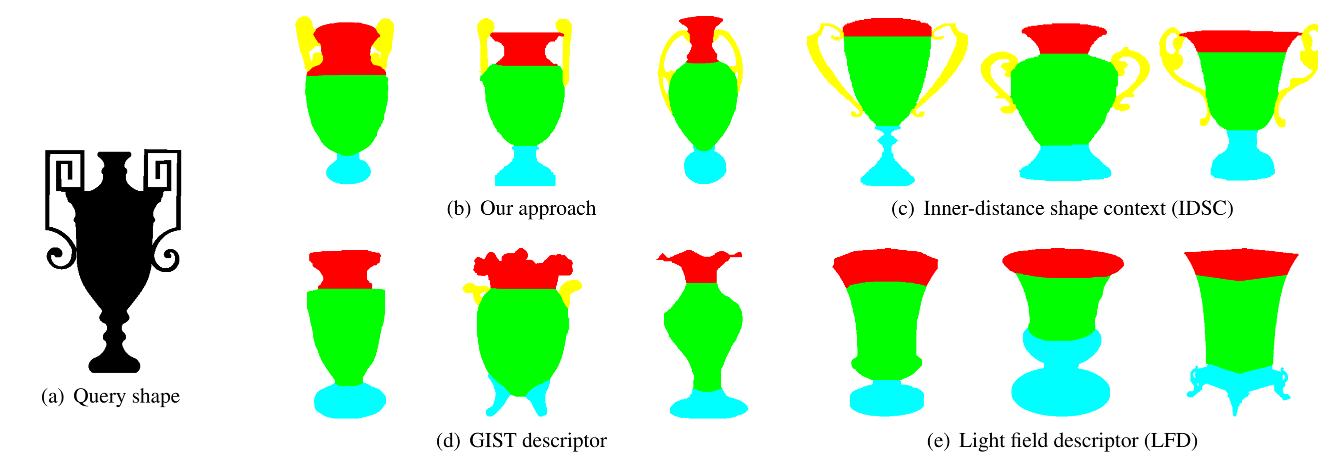

Figure 3: *Visual comparison for shapes retrieved by different approaches for a vase projection (a). Our approach (b) and IDSC (c) both return shapes that have similar overall topology. The remaining two approaches (c and d), are not able to retrieve shapes with similar topology.*

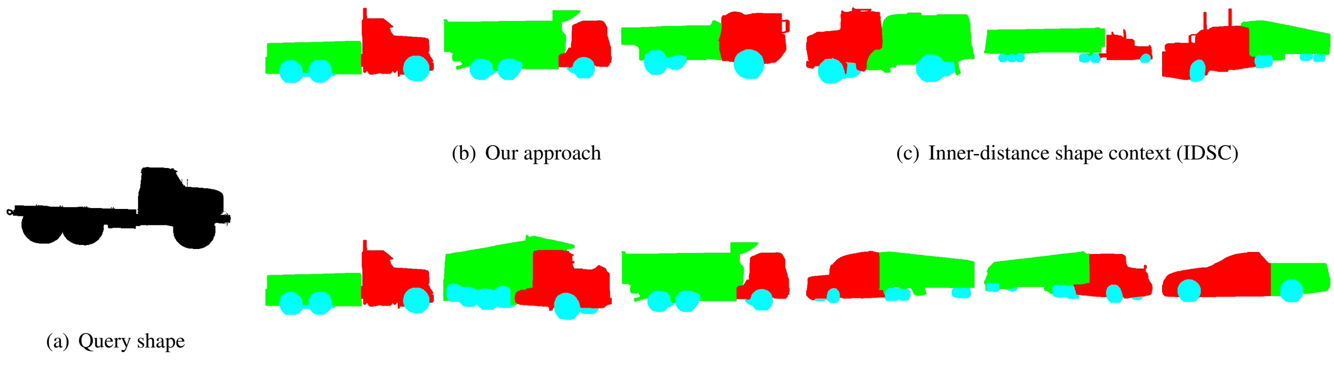

(a) Query shape

(b) Our approach

(c) Inner-distance shape context (IDSC)

(d) GIST descriptor

(e) Light field descriptor (LFD)

Figure 4: *Visual comparison for shapes retrieved by different approaches for a truck projection (a). Our approach (b) returns shapes that have similar overall topology and similar views as the projection. Some of the shapes retrieved by existing approaches have incorrect views.*

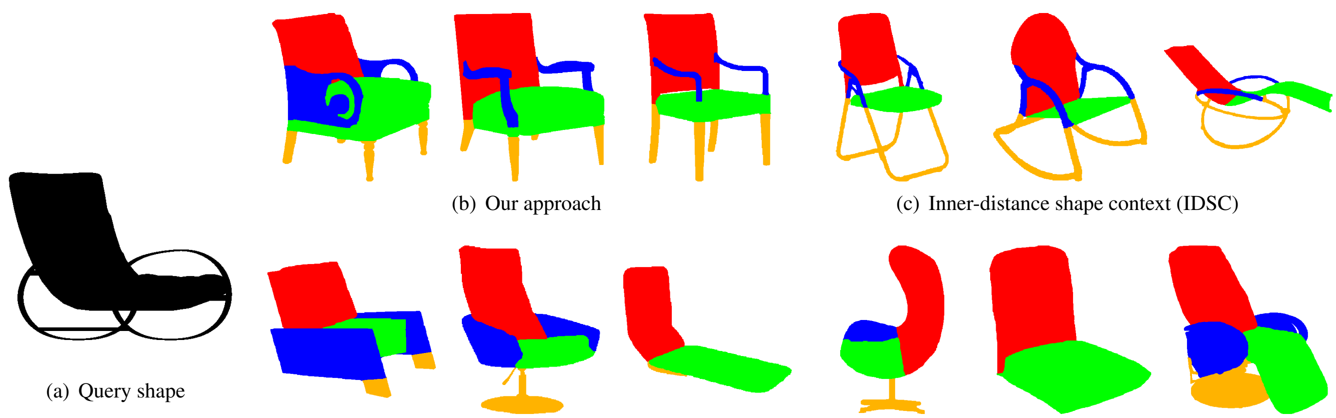

(a) Query shape

(b) Our approach

(c) Inner-distance shape context (IDSC)

(d) GIST descriptor

(e) Light field descriptor (LFD)

Figure 5: *An example showing our approach does not perform as well as IDSC. None of the chairs retrieved by our approach (b) has loops in their supporting structures, whereas the ones returned by IDSC (c) do. Please refer to Figure* **??***(c) for label transfer result.*

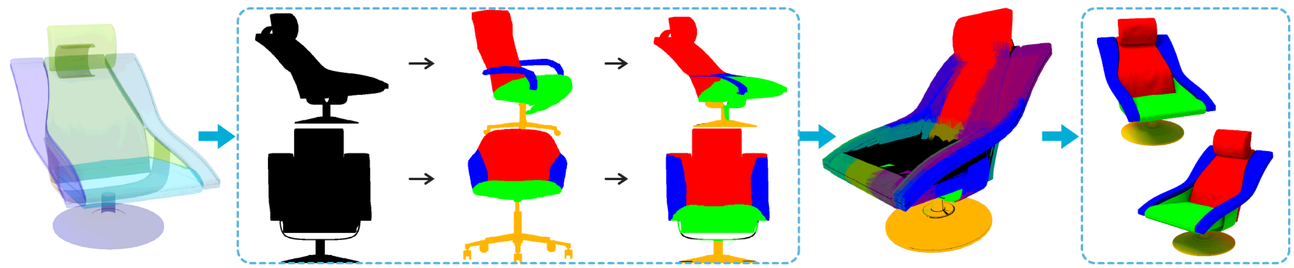

(a) Chair (10 pieces, 15674 triangles): Although no images in the database having similar supporting disc as the query shape, our approach properly labeled the two projections using chairs with wheels.

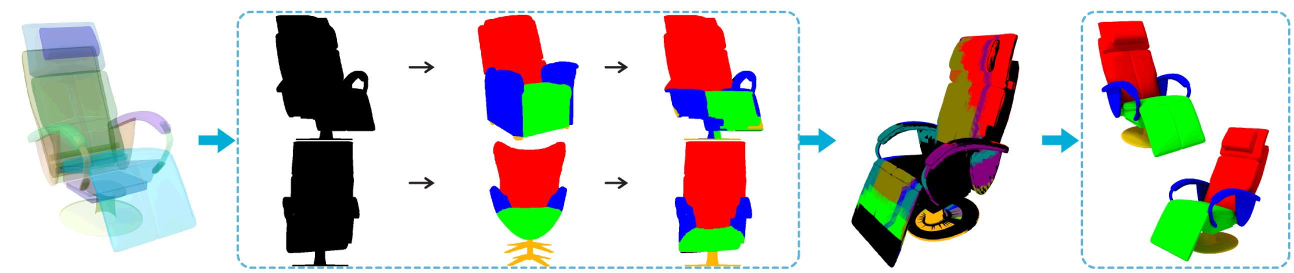

(b) Chair (18 pieces, 46540 triangles): Since the matched chairs have different topology for armrests as the query shape, some areas in the projections are incorrectly labeled. Nevertheless, combining multiple projections together generates a reasonable segmentation.

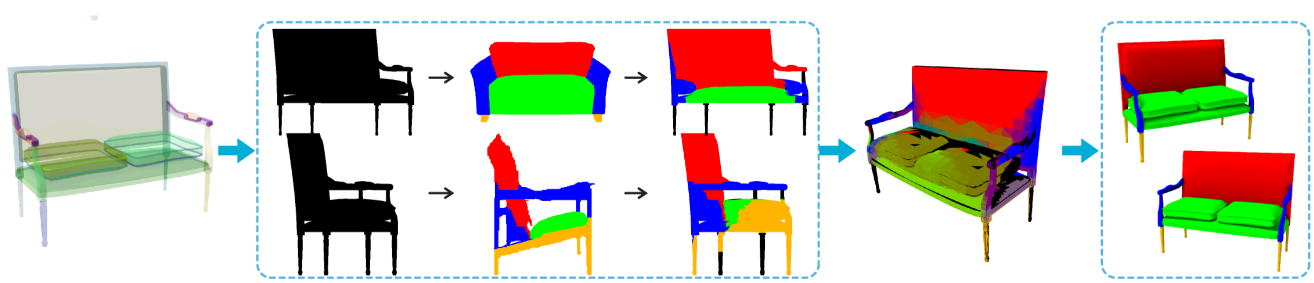

(c) Chair (26 pieces, 56004 triangles): Although there is no two-seat chair like the input shape in the database, our approach is able to infer proper labels from a two-seat sofa and a single seat chair.

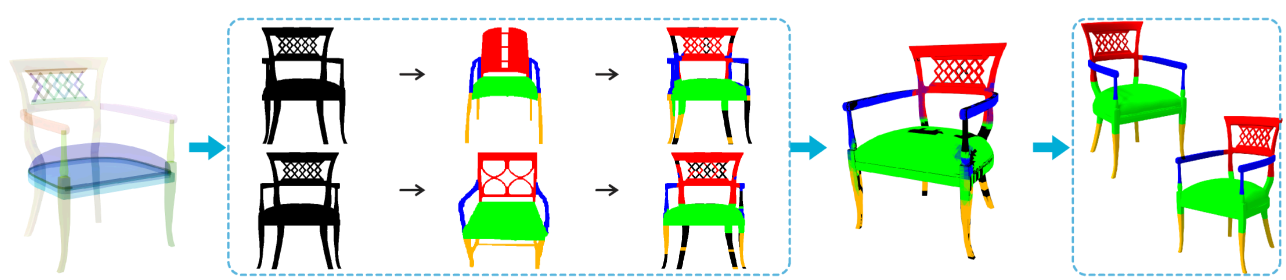

(d) Chair (19 pieces, 28664 triangles): When the input is a hollow back chair, our approach finds similar hollow back chairs while ignoring the fine-level topology differences, suggesting the BiSH distance measure is robust against small internal holes.

Figure 6: *Labeling results on various chair models downloaded from Trimble Warehouse.*

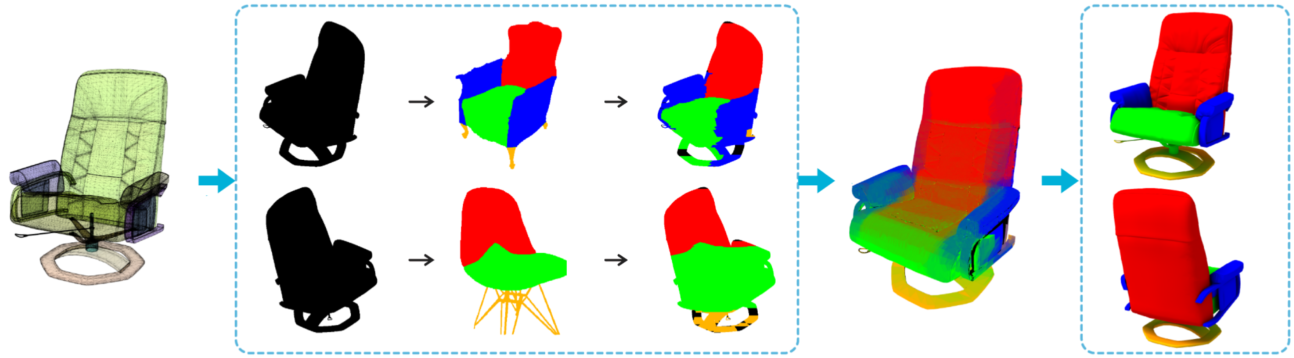

(a) Chair (16 pieces, 71582 triangles): Although not all matched labeled photos have armrests, combining them together can generate a reasonable segmentation.

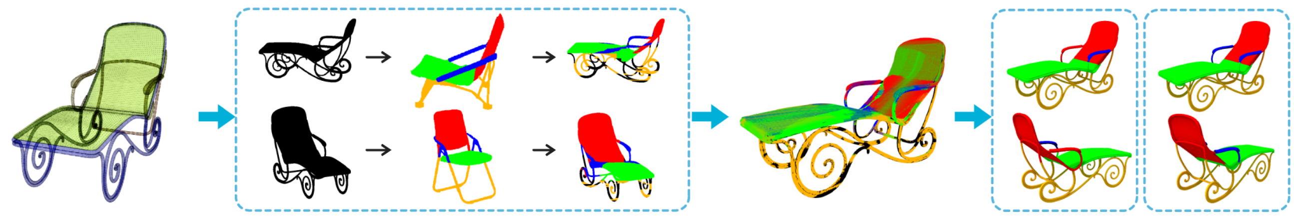

(b) Chair (5 pieces, 52624 triangles): In this input model, the back and the seat of the chair are grouped into a single piece. Hence, assuming triangles in the same piece having the same semantic label would give incorrect result. With the default settings, the left armrest is incorrectly labeled as back (red color). Setting $K_1 = 1$ and $K_2 = 1$, i.e., using only one photo for label transfer, gives a better result (last column).

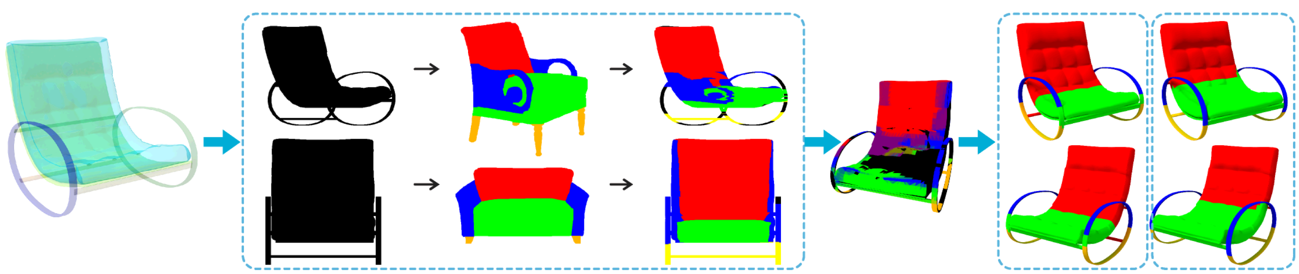

(c) Chair (8 pieces, 22844 triangles): This model contains a unique oval-shaped armrest-leg structure, which cannot be found in the real photo dataset. Nevertheless, a reasonable segmentation is obtained using the default parameters. Setting $K_1 = 5$ and $K_2 = 1$ further improves the labeling for the armrests (last column).

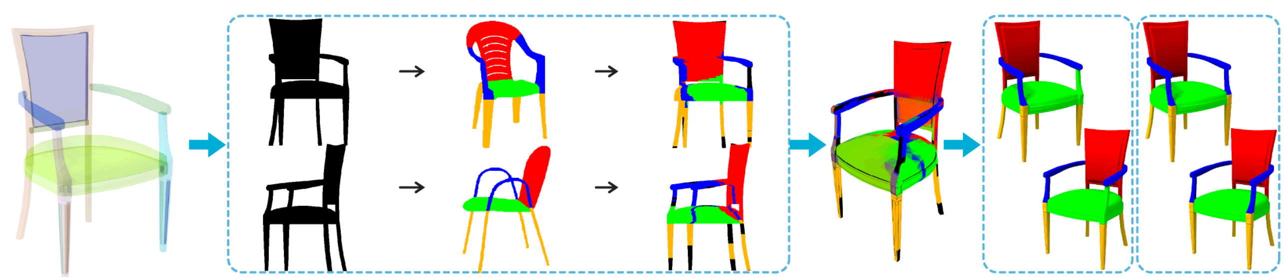

(d) Chair (21 pieces, 21036 triangles): Under the default setting, one of the armrests is partially mislabeled as seat (green color). Setting $K_1 = 3$ and $K_2 = 3$ gives a better result (last column).

Figure 7: *Labeling results on more chairs with different design styles.*

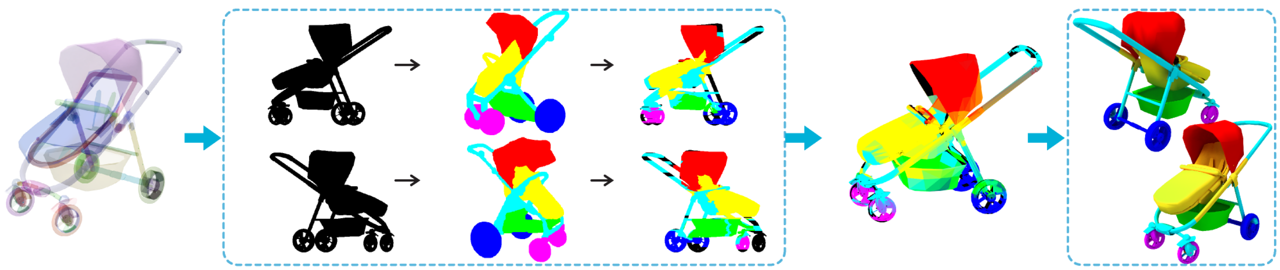

(a) Stroller (7 pieces, 17471 triangles): With six different labels, models in the stroller category are the most semantically complicated. Nevertheless, most parts in the input 3D shape are correctly labeled, except a small portion of the frame (cyan color) is incorrectly labeled as seat (yellow color).

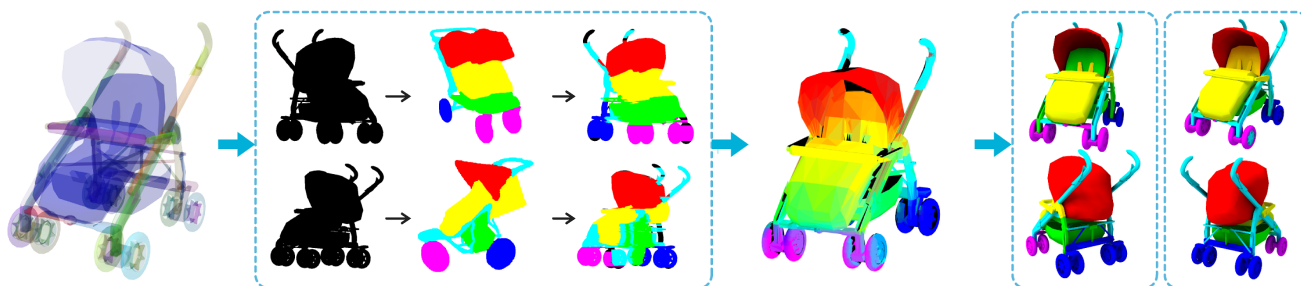

(b) Stroller (123 pieces, 32902 triangles): With the default settings, the baby seat (yellow color) is incorrectly labeled as green. Reducing the number of images retrieved for each projection (K_2) to 1 improves the result; see images in the last column.

Figure 8: *Labeling results on various stroller models downloaded from Trimble Warehouse.*

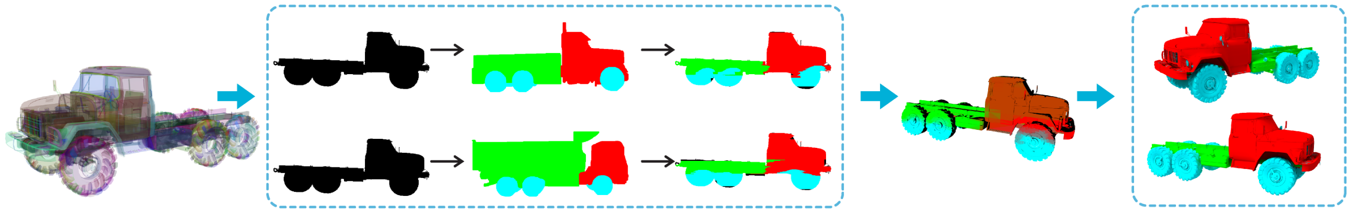

(a) Truck (1703 pieces, 41213 triangles): Although the retrieved images have different shapes for truck body as the query model, all major parts are correctly labeled. Note that some small parts (rear view mirror and front light) are mislabeled as wheels (cyan color).

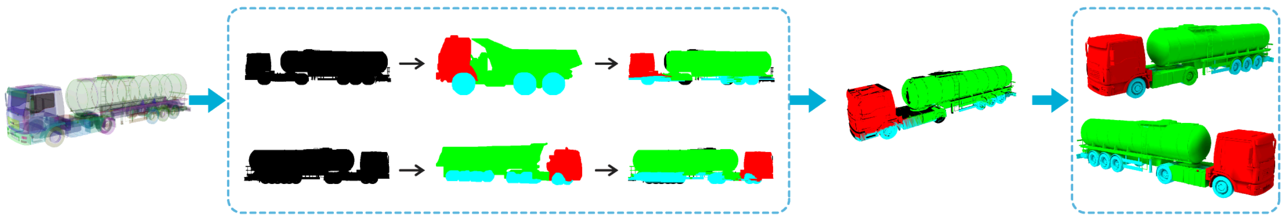

(b) Truck (392 pieces, 31432 triangles): Similar to the example above, most parts are correctly labeled.

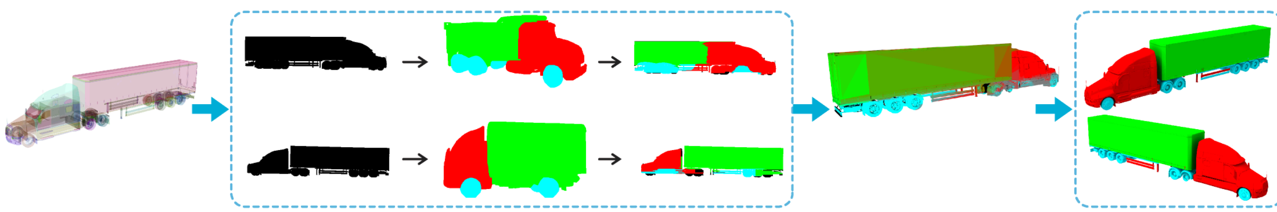

(c) Truck (852 pieces, 48264 triangles): Similar to the above example, our algorithm correctly labels most of the parts.

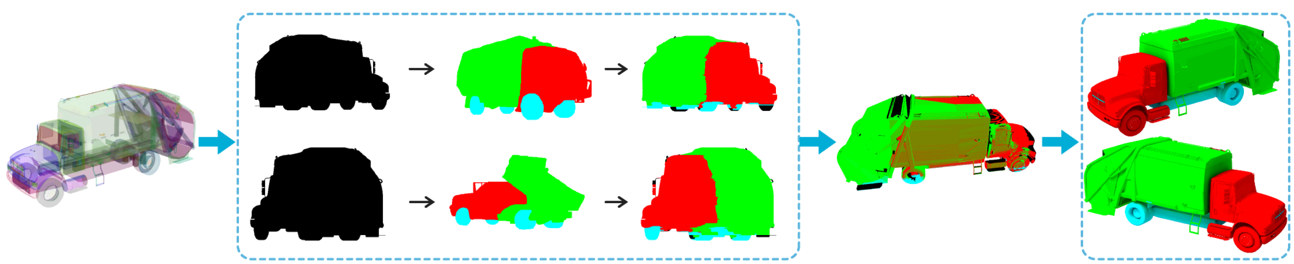

(d) Truck (493 pieces, 41213 triangles): A failure case, where the front wheels are labeled as head (red color) and part of the truck body is labeled as wheels (cyan color).

Figure 9: *Labeling results on various truck models downloaded from Trimble Warehouse.*

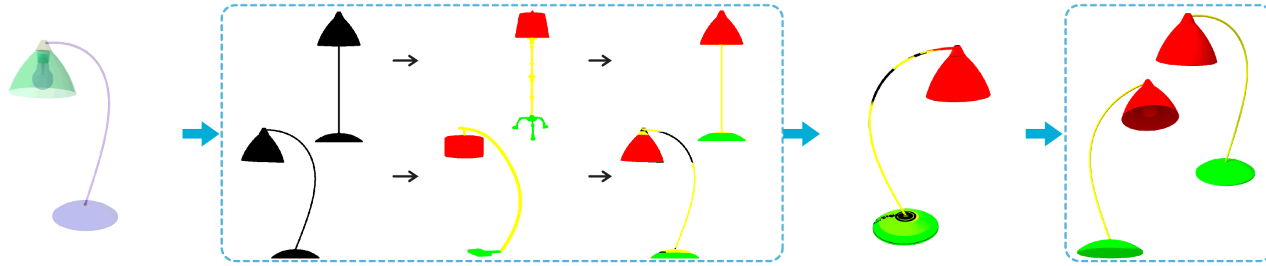

(a) Lamp (14 pieces, 76227 triangles): The images retrieved for different projections have quite different shapes. Our algorithm properly infers labeling information from these images.

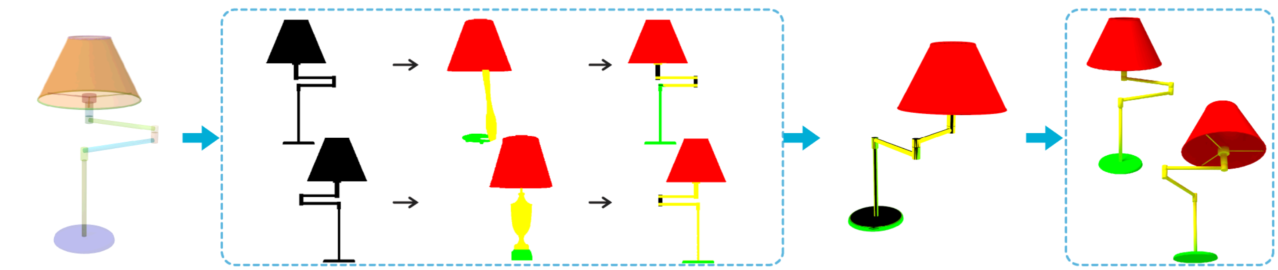

(b) Lamp (16 pieces, 4748 triangles): Similar with the above example, our algorithm accurately labels the input shape.

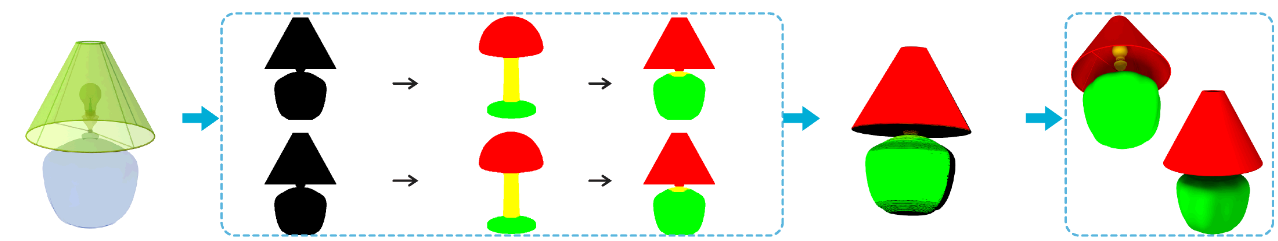

(c) Lamp (19 pieces, 165610 triangles): Due to the symmetrical design of the input lamp, the two projections have exact the same shape, resulting the same image being retrieved. Nevertheless, accurate labeling is obtained.

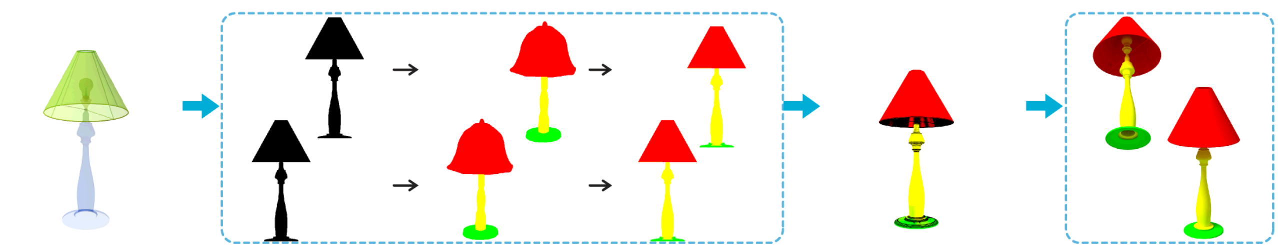

(d) Lamp (21 pieces, 140330 triangles): Similar with the above example, our algorithm achieves accurately labeling.

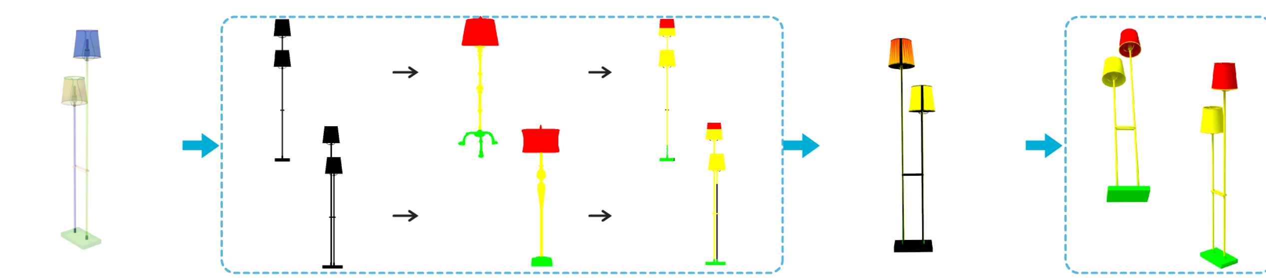

(e) Lamp (29 pieces, 13780 triangles): A failure case. Since none of the photos in labeled set has two lampshades, one of the lampshades is incorrectly labeled as stand (yellow color).

Figure 10: *Labeling results on various lamp models downloaded from Trimble Warehouse.*

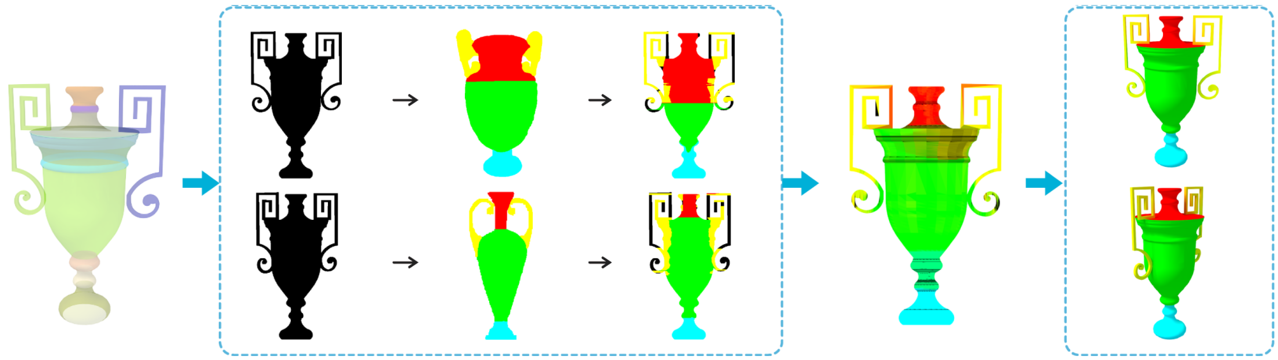

(a) Vase (14 pieces, 16801 triangles): The input vase has a complex handle design. Our algorithm proper labels it using photos of vases with much simpler handles.

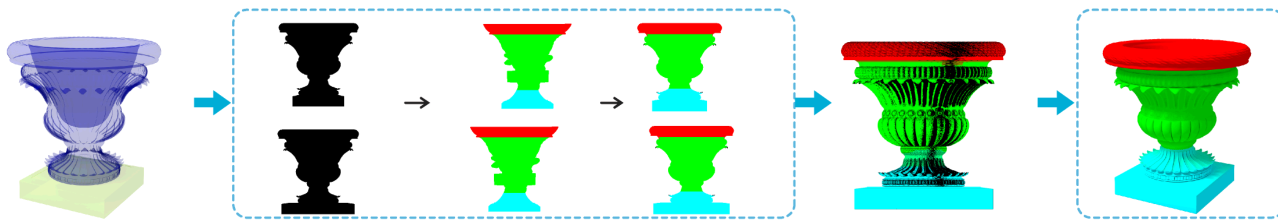

(b) Vase (2 pieces, 547308 triangles): Although two projections are matched to the same labeled photo, our algorithm properly labels the 3D shape.

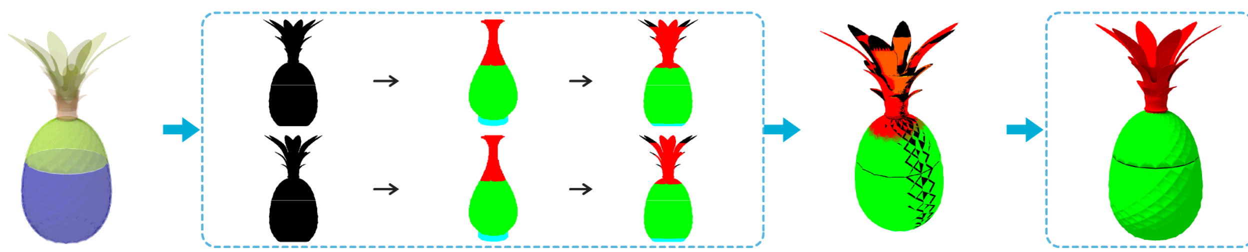

(c) Vase (5 pieces, 36716 triangles): Similar to the example above, the 3D shape can be correctly labeled.

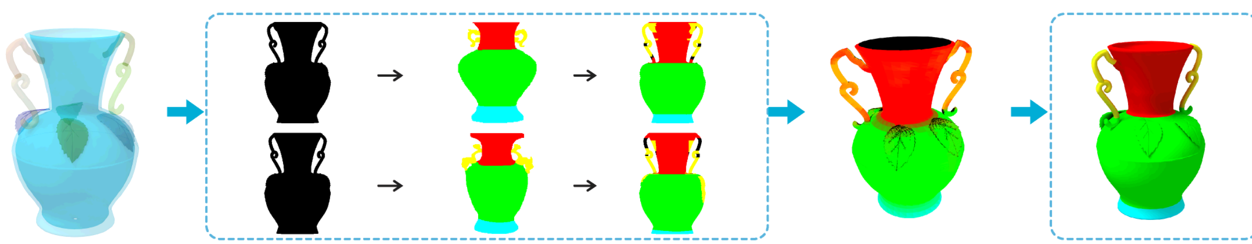

(d) Vase (9 pieces, 75952 triangles): Combing the labels transferred from two different shapes, our algorithm can achieve correctly labeling.

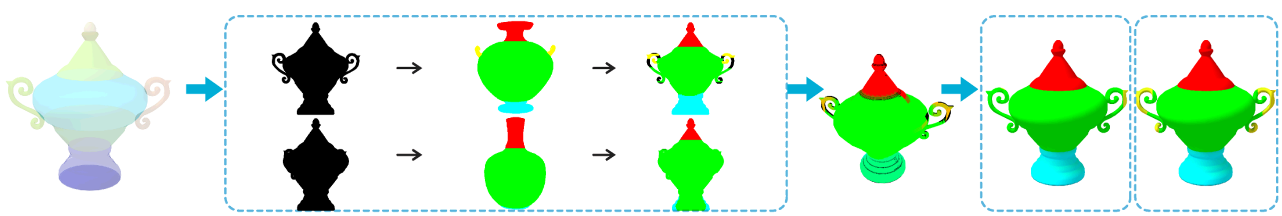

(e) Vase (8 pieces, 15402 triangles): Under default settings, the handle is incorrectly labeled as body (green color). Setting $K_1 = 5$ and $K_2 = 3$ gives a better result (last column).

Figure 11: *Labeling results on various vase models downloaded from Trimble Warehouse.*

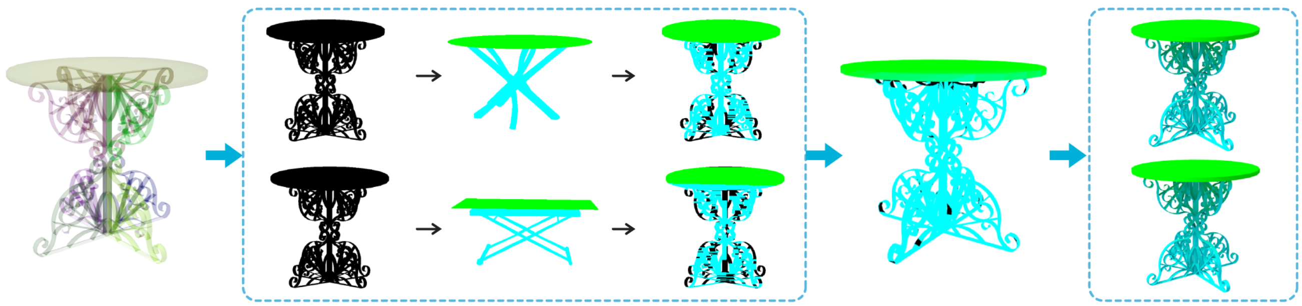

(a) Table (320 pieces, 8342 triangles): Since the proposed BiSH measure is robust against small internal holes, our algorithm properly retrieves and transfers labels from shapes that have similar overall topology, but much simpler in details.

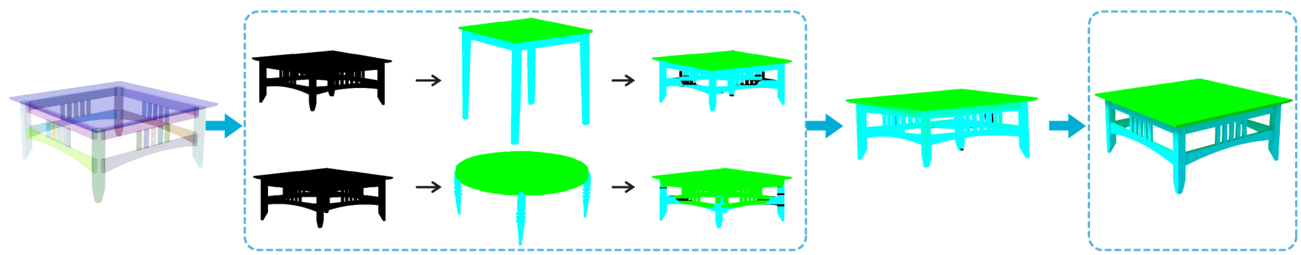

(b) Table (35 pieces, 1136 triangles): Matching between warp-aligned images allows our algorithm to properly infer lables from a table with much longer legs.

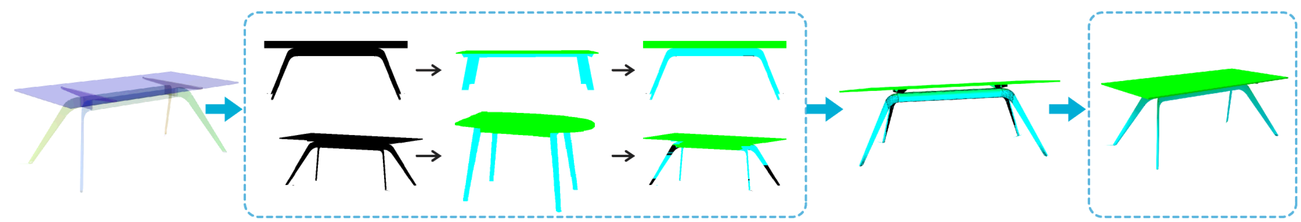

(c) Table (234 pieces, 31246 triangles): Similar to the example above, one of the tables retrieved is much higher than the input shape.

Figure 12: *Labeling results on various table models downloaded from Trimble Warehouse.*

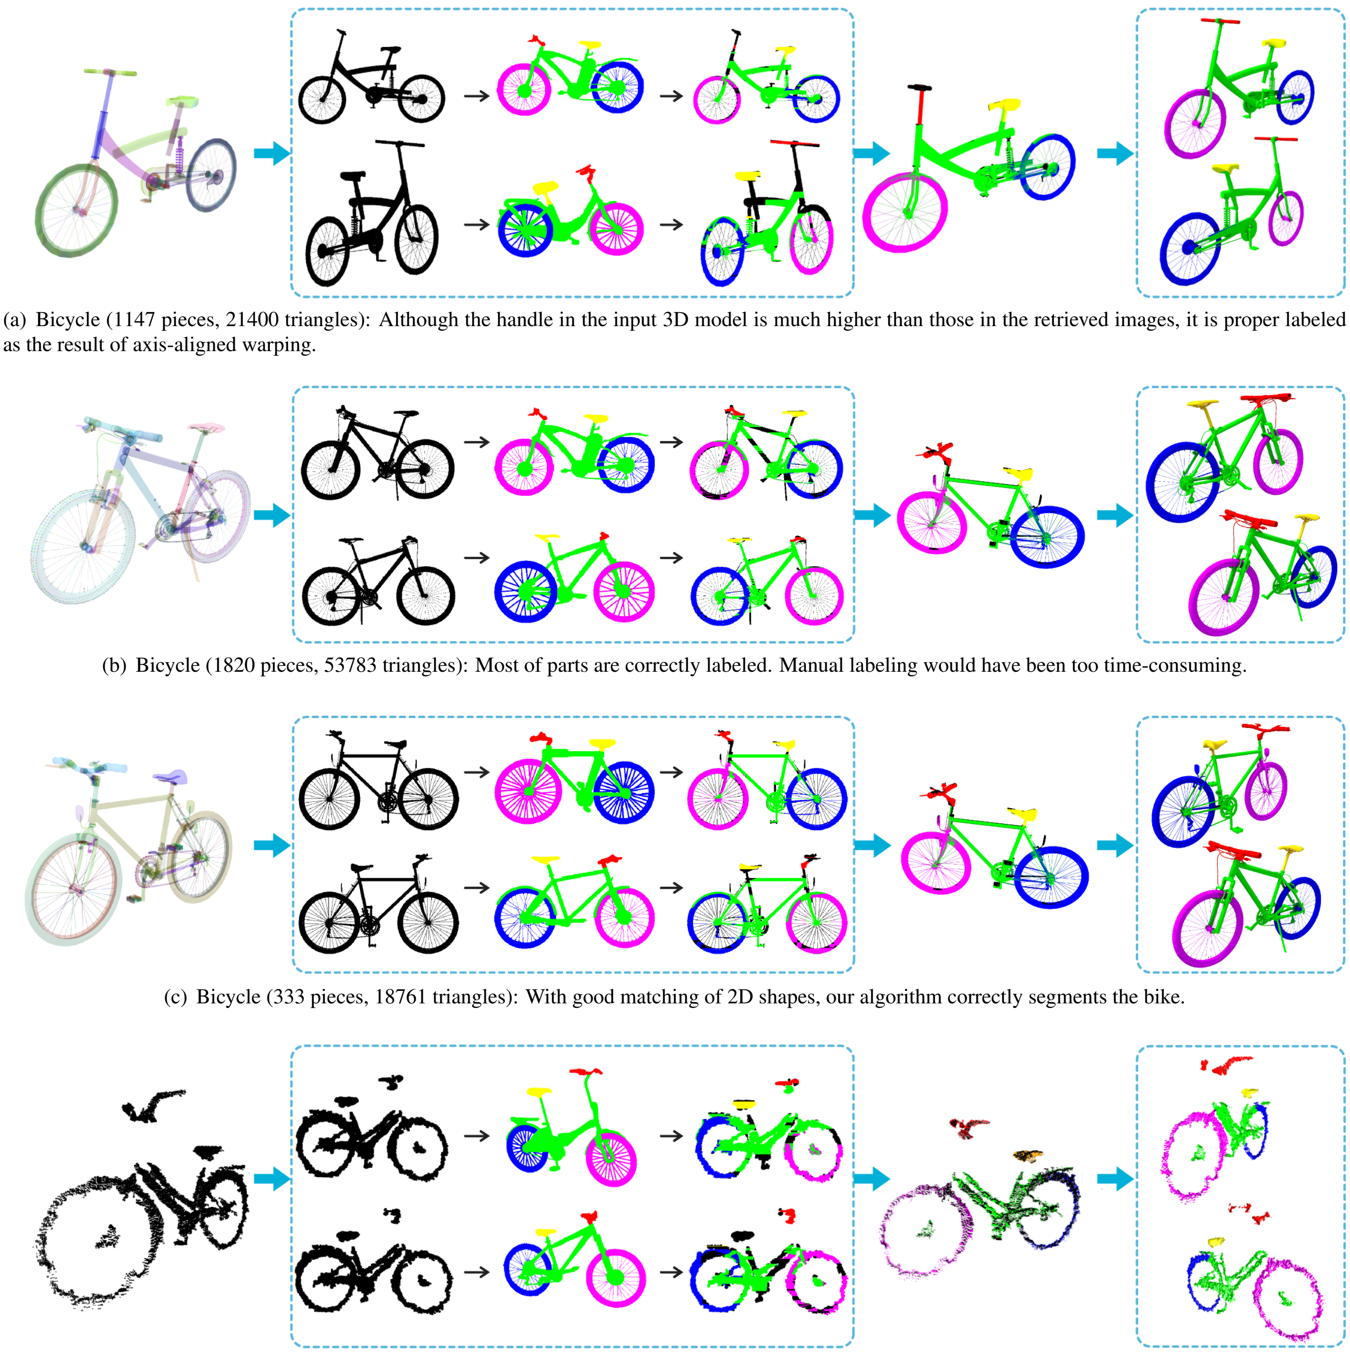

(a) Bicycle (1147 pieces, 21400 triangles): Although the handle in the input 3D model is much higher than those in the retrieved images, it is proper labeled as the result of axis-aligned warping.

(b) Bicycle (1820 pieces, 53783 triangles): Most of parts are correctly labeled. Manual labeling would have been too time-consuming.

(c) Bicycle (333 pieces, 18761 triangles): With good matching of 2D shapes, our algorithm correctly segments the bike.

(d) Raw scan of a bicycle (33182 points): Note that the scan is very sparse and highly incomplete.

Figure 13: *Labeling results on imperfect models and point cloud for bicycles.*

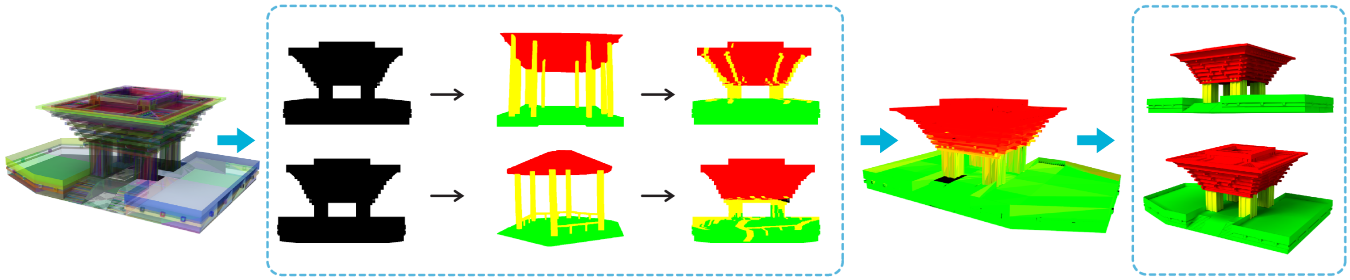

(a) Pavilion (3256 pieces, 88560 triangles): Note how the fine pieces are properly labeled in the final result. Manually label them can be very time-consuming.

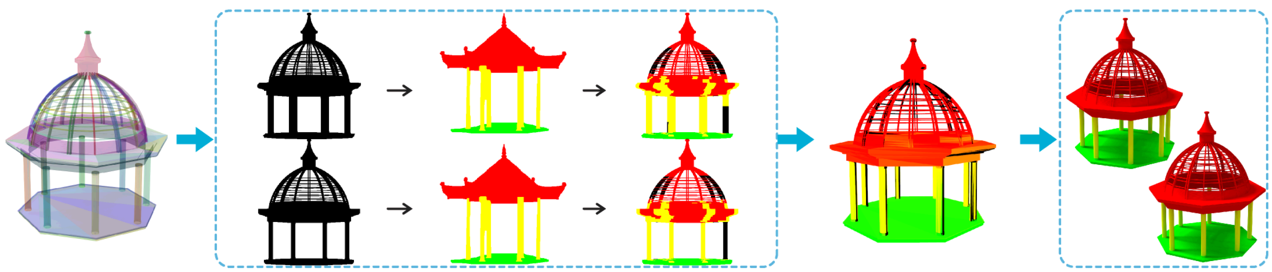

(b) Pavilion (719 pieces, 8625 triangles): Although two projections with different shapes match to the same labeled photo, the pavilion is still correctly segmented.

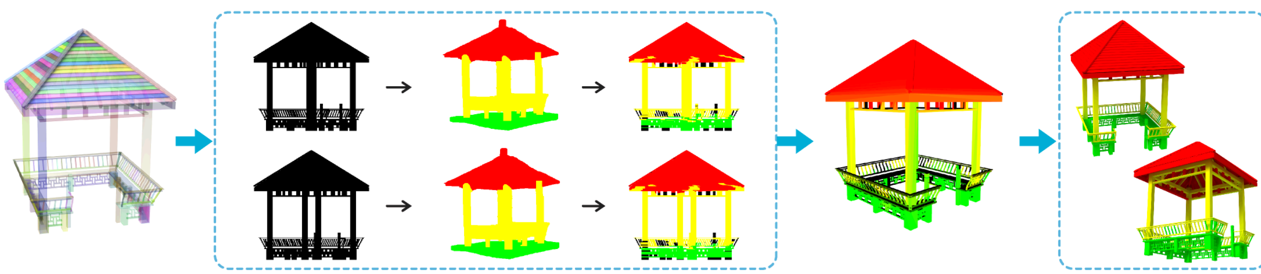

(c) Pavilion (903 pieces, 19495 triangles): A failure case. The input model does not contain a ground plane as the base. As a result, parts of the supporting polars (yellow color) are incorrectly labeled as base (green color).

Figure 14: *Labeling results on various pavillion models downloaded from Trimble Warehouse.*

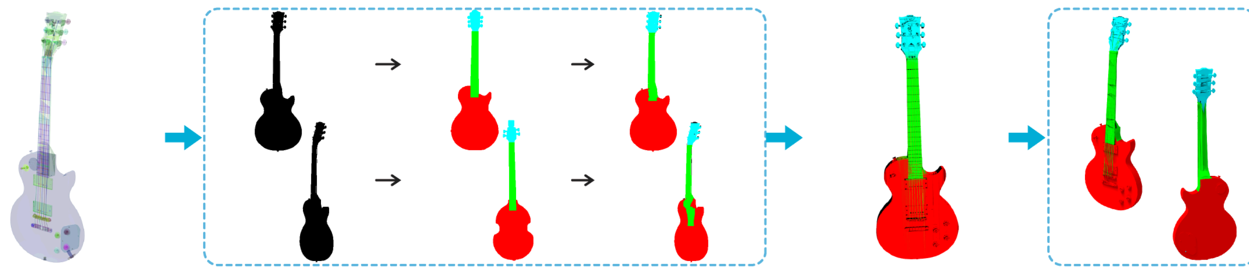

(a) Guitar (326 pieces, 76227 triangles): For shapes with simple topology such as guitars, our algorithm can properly label different parts using a small labeled set of only 20 labeled images.

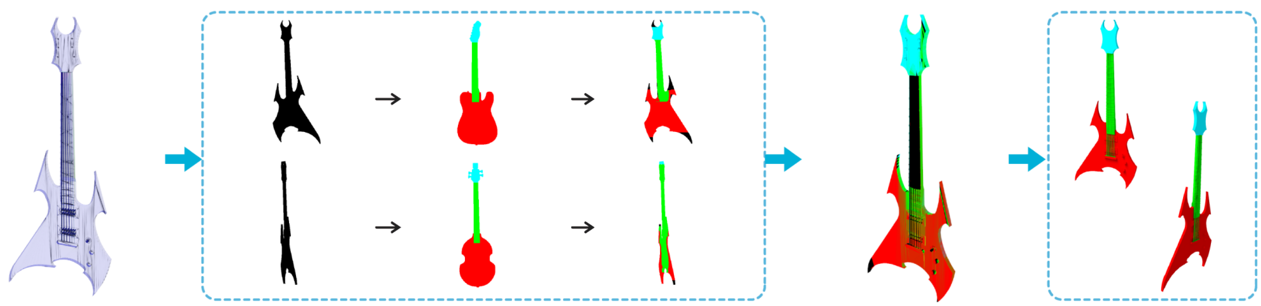

(b) Guitar (87 pieces, 6287 triangles): Most parts are properly labeled for this guitar with unconventional shape.

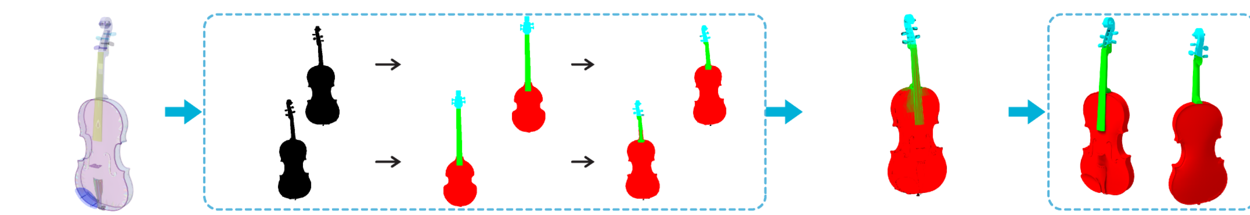

(c) Violin (3778 pieces, 32786 triangles): Segmenting this violin model using labeled guitar images gives mostly correct result as well.

Figure 15: *Labeling results on guitar and violin models downloaded from Trimble Warehouse.*

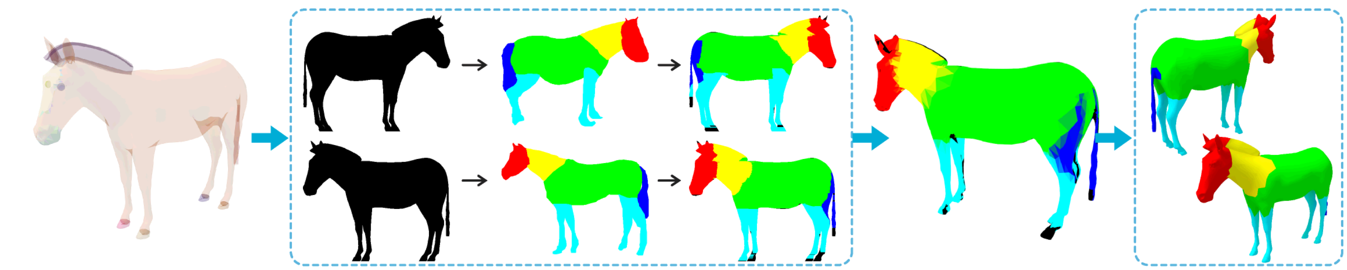

(a) Fourleg (112 pieces, 8741 triangles): Although designed for handling rigid objects, our algorithm can also label articulated objects as long as shapes with similar articulation settings exist in the labeled dataset.

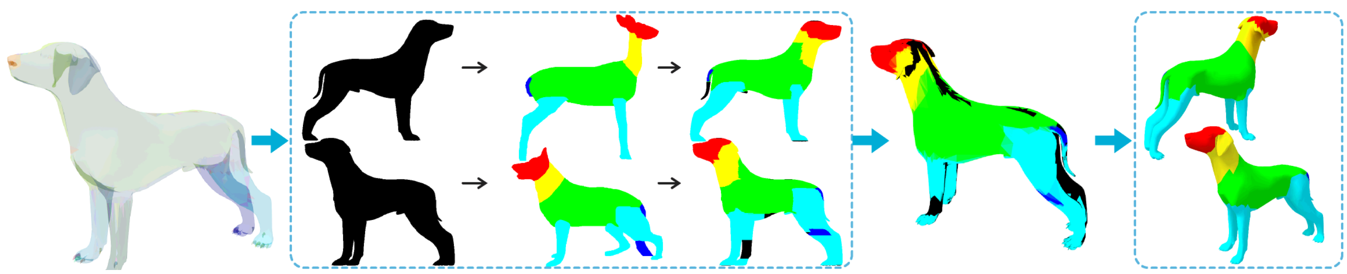

(b) Fourleg (522 pieces, 20746 triangles): Due to the deformation in the tail, only a part of it is correctly labeled.

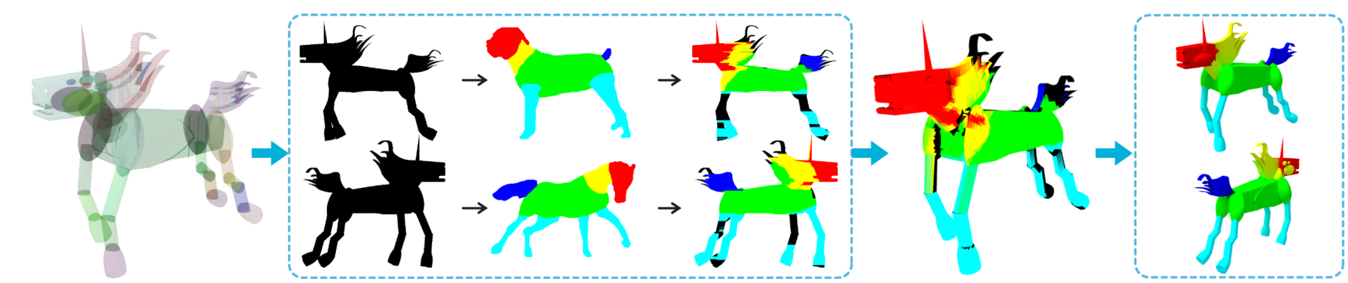

(c) Fourleg (908 pieces, 16186 triangles): By inferring knowledge from photos of real animals, our algorithm can properly label a very robotic looking animal model.

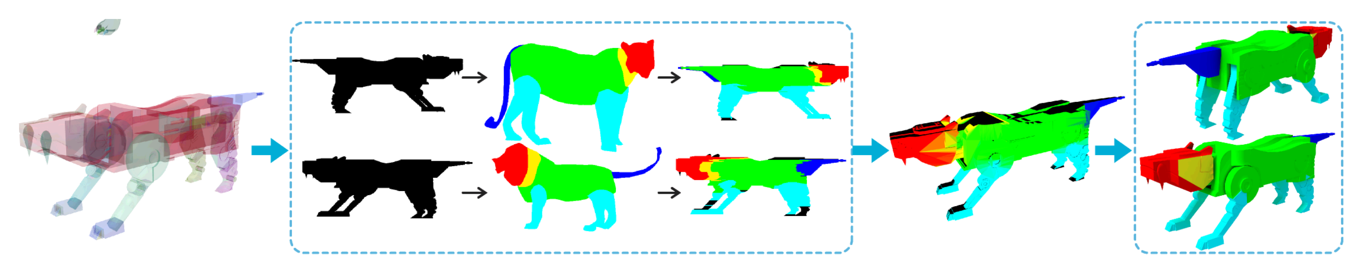

(d) Fourleg (928 pieces, 12497 triangles): A failure case. Since all animals retrieved have very short necks, the neck label (yellow color) is not properly transferred to the input model.

Figure 16: *Labeling results on four-leg animal and robot models downloaded from Trimble Warehouse.*

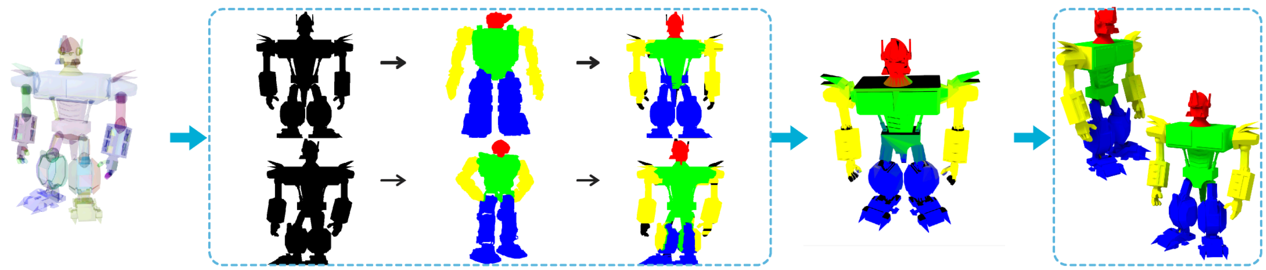

(a) Robot (1168 pieces, 25739 triangles): The input robot model is properly labeled using photos of robots with similar poses.

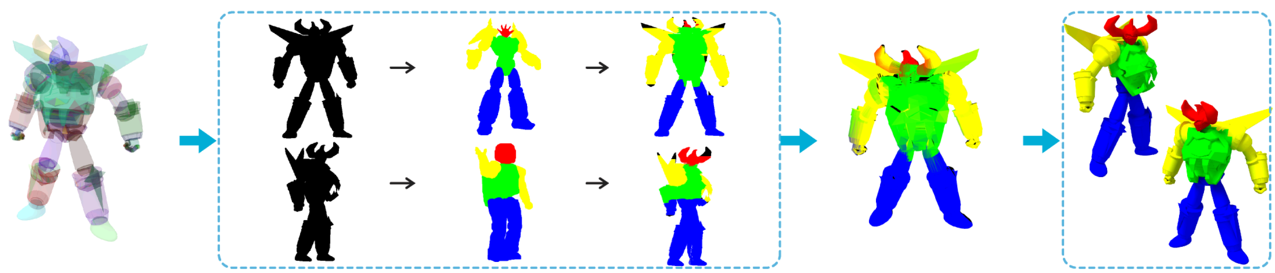

(b) Robot (1696 pieces, 76227 triangles): In the bottom projection, the wings of the input robot are incorrectly labeled as arm (yellow color), whereas an arm is incorrectly labeled as body (green) and leg (blue). However, the final labeling result is correct.

Figure 17: *Labeling results on various robot models downloaded from Trimble Warehouse.*

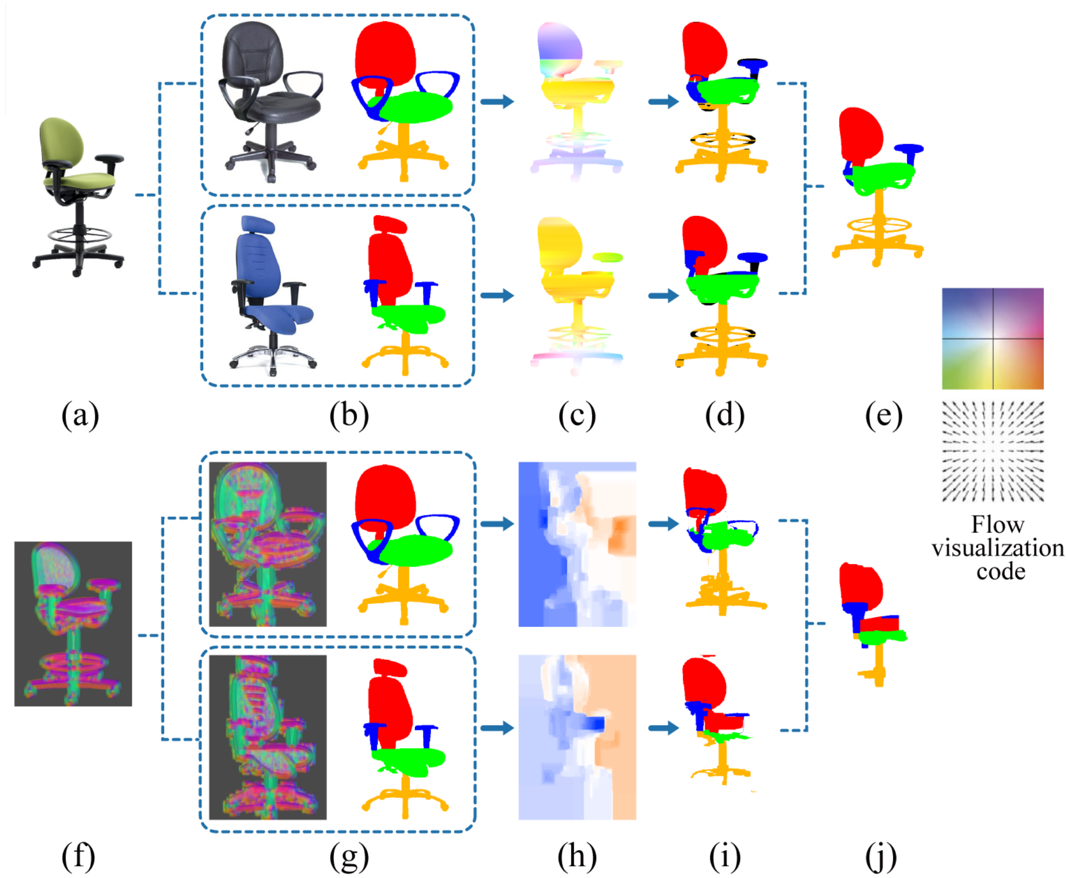

Figure 18: *Labeling of a 2D chair image. This is the same figure as Figure 15 in the paper, but with bigger size for better visibility. The top row shows the result of our approach, whereas the bottom row shows the one obtained using [Liu et al. 2011a].*

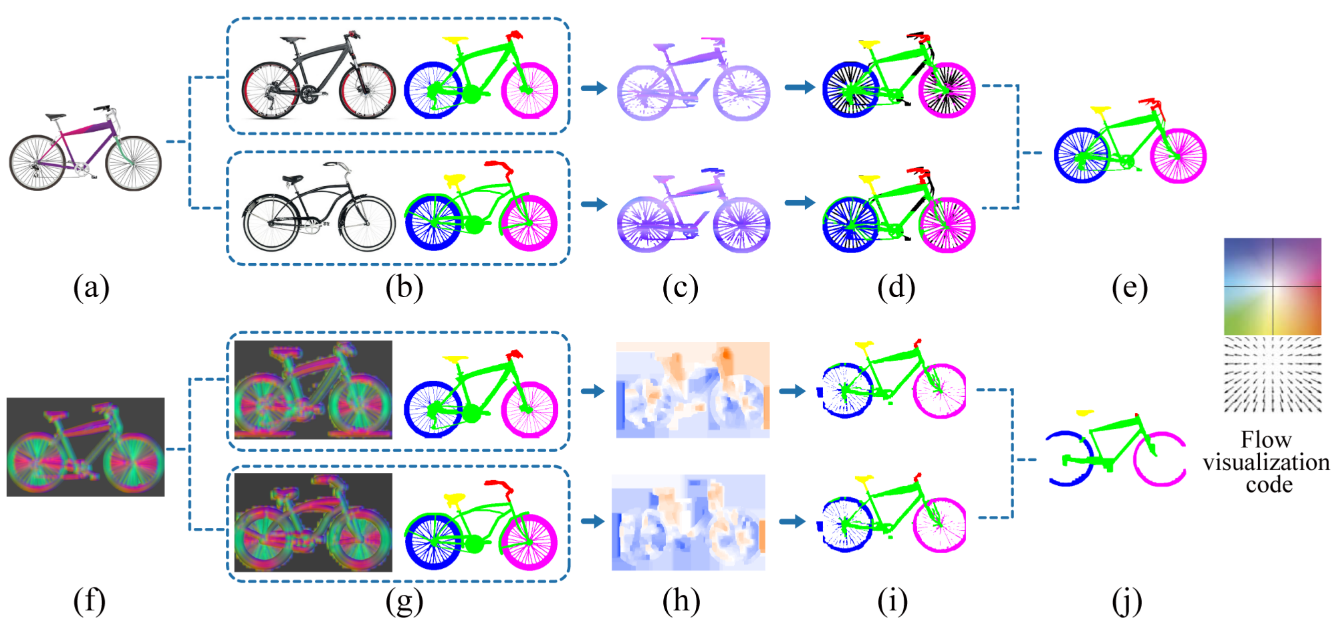

Figure 19: *Labeling of a 2D bicycle image. The layout of subfigures are the same as Figure* **??**.