# Enhanced Back-Projection of Vision Features for 3D Symmetry Detection

Isaac Aguirre, Ivan Sipiran
Department of Computer Science, University of Chile
{iaguirre, isipiran}@dcc.uchile.cl

## Supplementary Material

This supplementary material provides additional details and visualizations that complement the main paper. We include: (i) comparisons of viewpoint sampling strategies, (ii) detailed timing analysis for plane and axis detection, (iii) the effect of group size on axial symmetry, (iv) additional robustness experiments under noise and part removal, and (v) qualitative results illustrating back-projected features and detected symmetries using PCA-colored visualizations. These materials aim to give readers a deeper understanding of our method's performance and behavior.

## A. Viewpoint Sampling

Standard latitude–longitude sampling (RE) places viewpoints by uniformly stepping azimuth and elevation. This creates dense clusters near the poles even when using the same number of views $V$. Fibonacci sphere sampling (FI) places viewpoints using increments of the golden angle, producing a near-uniform distribution on the sphere. This leads to more even coverage with the same $V$. Figure 1 compares both view sampling methods. Both samplings (RE and FI) are provided in our public code release linked in the abstract.

## B. Detailed Timing Analysis

Tables 1 and 2 provide a stage-wise breakdown of computation times for our plane and axis detection algorithms, where Feature Extraction and Plane Detection are the average per object.

## C. Effect of Group Size $L$ on Axial Symmetry

Table 3 reports detailed results for axial symmetry across group sizes $L = 10, 20, 30, 40, 50$, complementing the summary provided in Results section of the main paper. As noted there, performance consistently improves with larger groups, with $L = 50$ yielding the best overall results.

## D. Additional Robustness Results (Noise and Part Removal)

Table 4 presents the full version of Table 2 of the main paper. In addition, Figure 2 illustrates an example of robustness for planes; results for axes are omitted since their robustness is comparatively weaker.

## E. Qualitative Results

### Visualization of Back-Projected Features and Symmetry Detection

To provide qualitative insight into our method, we show the back-projected features and detected symmetry structures for several objects. All visualizations are color-coded using PCA projections of the features to illustrate feature similarity across symmetric points.

**Planes** Figures 3 to 5 show six representative objects for plane detection. For each object, we present two visualizations: the left image displays the colored point cloud of back-projected features using our best configuration for planes (4R-FI, FM10K, 114 views), while the right image shows the corresponding mesh (4R-FI, RM, 114 views) with the detected plane overlaid. The mesh is included because the plane is not easily visible in the point cloud alone, and it improves interpretability while still showing feature consistency across symmetric regions. Similar colors indicate feature similarity along symmetric points.

**Axes** Figure Fig. 6 shows six representative objects for axis detection. Each image shows the colored mesh using the best configuration for axes (4R-FI, RM, 114 views, L50) with the detected axes drawn as blue lines passing through the center of the object. The coloring corresponds to PCA-projected back-projected features, demonstrating feature similarity along symmetric regions.

Figure 1. Viewpoint distributions on the sphere. Left: standard $RE$ sampling concentrates views near the poles. Right: Fibonacci $FI$ sampling yields a more even spread with the identical $V = 114$.

| Configuration | Views | Total Time (s) | Avg per Object (s) | Feature Extraction (s) | Plane Detection (s) |
|---|---|---|---|---|---|
| RM-1R-RE | 6 | 564.6 | 6.643 | 0.531 | 6.058 |
| | 114 | 1123.7 | 13.220 | 6.884 | 6.269 |
| RM-4R-FI | 6 | 608.9 | 7.164 | 0.721 | 6.393 |
| | 114 | 1778.0 | 20.918 | 13.309 | 7.554 |
| PC10K-1R-RE | 6 | 536.8 | 6.315 | 0.101 | 6.167 |
| | 114 | 685.5 | 8.065 | 1.333 | 6.687 |
| PC10K-4R-FI | 6 | 559.6 | 6.584 | 0.244 | 6.293 |
| | 114 | 1055.9 | 12.423 | 3.949 | 8.424 |
| FM10K-1R-RE | 6 | 519.0 | 6.106 | 0.562 | 5.498 |
| | 114 | 1061.9 | 12.493 | 6.803 | 5.625 |
| FM10K-4R-FI | 6 | 537.1 | 6.318 | 0.739 | 5.530 |
| | 114 | 1702.7 | 20.032 | 13.320 | 6.658 |

Table 1. Stage-wise timing for plane detection over 85 objects (average size: 4,424 points/vertices).

| Configuration | Views | Total Time (s) | Avg per Object (s) | Feature Extraction (s) | Axis Detection (s) |
|---|---|---|---|---|---|
| RM-1R-RE-L50 | 6 | 6564.9 | 77.235 | 0.487 | 76.638 |
| | 114 | 5471.5 | 64.371 | 5.520 | 58.756 |
| RM-4R-FI-L50 | 6 | 3401.6 | 40.018 | 0.707 | 39.222 |
| | 114 | 4080.8 | 48.009 | 10.166 | 37.748 |

Table 2. Stage-wise timing for axis detection over 85 objects (average size: 16,177 points/vertices).

| Configuration | Views | L size | F-score | SDE($\times 10^{-4}$) |
|---|---|---|---|---|
| FM10K-4R-FI | 6 | 10 | 0.037 | 702.0 |
| Feature-Mesh | | 20 | 0.053 | 728.9 |
| Sampling | | 30 | 0.065 | 638.9 |
| 10K points | | 40 | 0.085 | 575.9 |
| (Enhanced) | | 50 | 0.096 | 521.8 |
| | 42 | 10 | 0.067 | 617.5 |
| | | 20 | 0.116 | 501.6 |
| | | 30 | 0.126 | 366.6 |
| | | 40 | 0.144 | 311.1 |
| | | 50 | 0.185 | 314.7 |
| | 114 | 10 | 0.085 | 468.4 |
| | | 20 | 0.137 | 309.5 |
| | | 30 | 0.172 | 246.8 |
| | | 40 | 0.196 | 199.9 |
| | | 50 | 0.221 | 229.4 |
| PC10K-4R-FI | 6 | 10 | 0.016 | 2728.4 |
| Point Clouds | | 20 | 0.017 | 2539.5 |
| 10K points | | 30 | - | 1909.3 |
| (Enhanced) | | 40 | 0.027 | 1642.2 |
| | | 50 | 0.040 | 1338.6 |
| | 42 | 10 | 0.087 | 637.1 |
| | | 20 | 0.133 | 503.0 |
| | | 30 | 0.174 | 513.5 |
| | | 40 | 0.184 | 580.2 |
| | | 50 | 0.167 | 451.0 |
| | 114 | 10 | 0.157 | 417.4 |
| | | 20 | 0.192 | 454.5 |
| | | 30 | 0.211 | 471.7 |
| | | 40 | 0.241 | 504.4 |
| | | 50 | 0.243 | 501.1 |
| RM-4R-FI | 6 | 10 | 0.381 | 430.7 |
| Raw-Mesh | | 20 | 0.497 | 397.4 |
| (Enhanced) | | 30 | 0.559 | 380.3 |
| | | 40 | 0.579 | 365.8 |
| | | 50 | 0.598 | 358.0 |
| | 42 | 10 | 0.583 | 275.0 |
| | | 20 | 0.686 | 260.2 |
| | | 30 | 0.722 | 252.2 |
| | | 40 | 0.747 | 252.9 |
| | | 50 | 0.763 | 249.0 |
| | 114 | 10 | 0.649 | 229.5 |
| | | 20 | 0.736 | 229.9 |
| | | 30 | 0.764 | 226.2 |
| | | 40 | 0.781 | 239.8 |
| | | 50 | 0.803 | 231.9 |

Table 3. Performance of Raw-Mesh (RM), Feature-Mesh Sampling (FM), and Point Clouds (PC) (FM and PC sampled with 10K points) on the axial symmetry algorithm across group sizes $L = 10, 20, 30, 40, 50$ and numbers of views (6, 42, 114) using the enhanced 4-rotation Fibonacci sampling (4R-FI). RM achieves the highest performance, and overall accuracy increases with both group size and number of views, reaching a maximum at $L = 50$ and 114 views.

| Dataset | Configuration | F-score | | | | | SDE($\times 10^{-4}$) | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | S | N1% | N3% | N5% | PR20% | S | N1% | N3% | N5% | PR20% |
| Planes | RM-4R-FI | 0.887 | 0.865 | 0.724 | 0.570 | 0.464 | 1.690 | 1.287 | 1.000 | 0.692 | 2.946 |
| | PC10K-4R-FI | 0.803 | 0.719 | 0.359 | 0.095 | 0.358 | 9.363 | 1.068 | 0.276 | 0.104 | 11.30 |
| | FM10K-4R-FI | 0.912 | 0.906 | 0.888 | 0.838 | 0.720 | 0.282 | 0.343 | 0.477 | 0.479 | 0.740 |
| Axes | FM10K-4R-FI-L50 | 0.221 | 0.125 | 0.095 | 0.076 | 0.147 | 229.4 | 165.1 | 115.3 | 83.04 | 239.0 |
| | PC10K-4R-FI-L50 | 0.243 | 0.074 | 0.055 | 0.050 | 0.152 | 501.1 | 849.7 | 474.9 | 261.1 | 857.1 |
| | RM-4R-FI-L50 | 0.803 | 0.301 | 0.223 | 0.182 | 0.581 | 231.9 | 197.0 | 143.0 | 117.0 | 238.9 |

Table 4. Performance of Raw-Mesh (RM), Feature-Mesh Sampling (FM), and Point Clouds (PC) (FM and PC sampled with 10K points) using the enhanced 4-rotation Fibonacci sampling (4R-FI) on planes and axes under different evaluation settings. Higher F-score indicates better detection accuracy, while lower SDE indicates better geometric alignment. All results use 114 views.

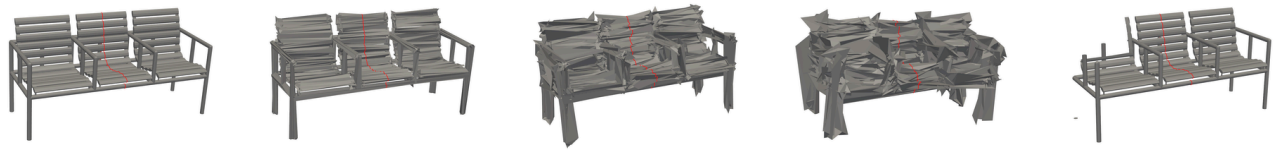

(a) (b) (c) (d) (e)

Figure 2. Example illustrating the robustness of the plane detection algorithm under its best configuration (FM10K-4R-FI): (a) original object, (b–d) with 1%, 3%, and 5% noise, and (e) with 20% part removal.

Figure 3. Qualitative results for planar symmetry detection (Part 1). Left: colored point cloud showing back-projected features. Right: mesh with detected plane overlay; the mesh highlights the plane since it is not clearly visible in the point cloud alone. Colors indicate feature similarity across symmetric points.

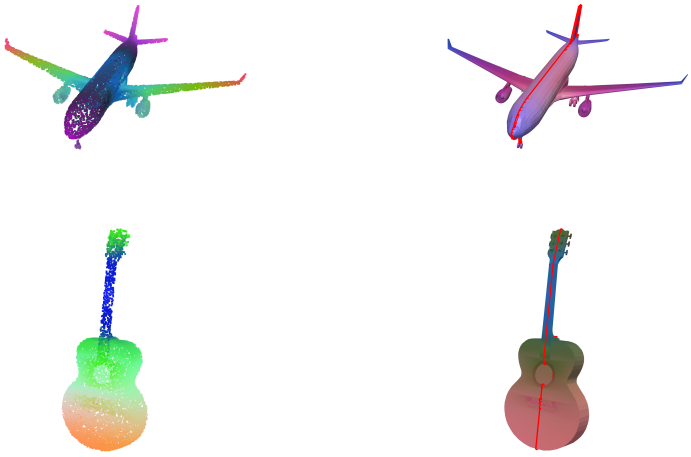

Figure 4. Qualitative results for planar symmetry detection (Part 2). Left: colored point cloud showing back-projected features. Right: mesh with detected plane overlay; the mesh highlights the plane since it is not clearly visible in the point cloud alone. Colors indicate feature similarity across symmetric points.

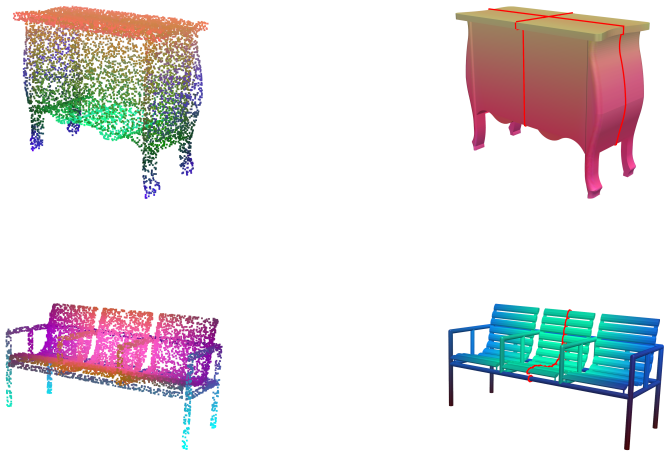

Figure 5. Qualitative results for planar symmetry detection (Part 3). Left: colored point cloud showing back-projected features. Right: mesh with detected plane overlay; the mesh highlights the plane since it is not clearly visible in the point cloud alone. Colors indicate feature similarity across symmetric points.

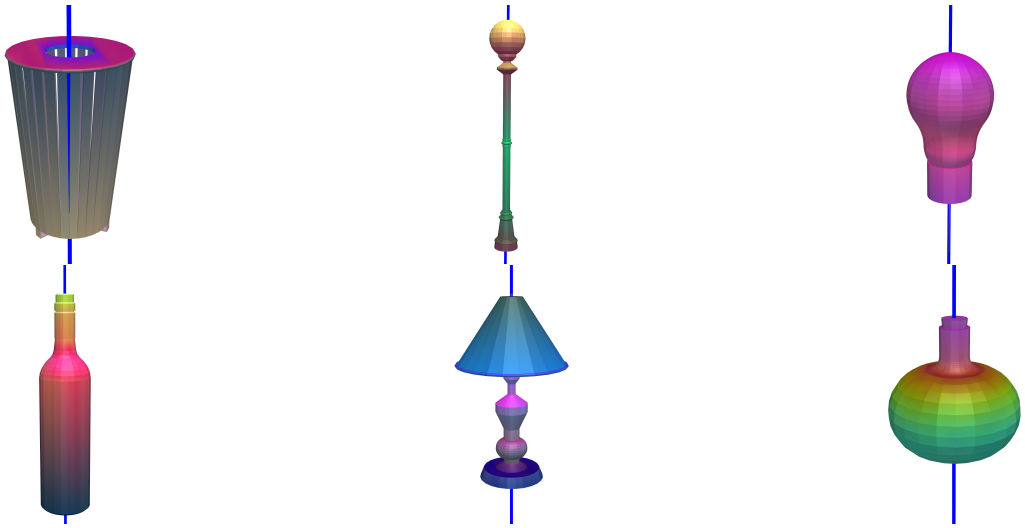

Figure 6. Qualitative results for rotational symmetry detection. Each colored mesh shows detected axes overlaid in blue. Colors indicate PCA-projected back-projected features, illustrating symmetry invariance.