A Appendix

A.1 Joint vs per category training

All methods presented in Tables 2 and 3 in the main paper, including *eigenSDF*, were trained jointly on all 13 ShapeNet categories. However, since methods presented in Tables 1, 4 and Figures 2, 4, were originally trained per category, we have retrained our *eigenSDF* model accordingly. We have observed, however, that final results did not differ substantially. The difference in average performance was smaller than 0.02 on all 3 metrics, i.e. Intersection over Union, Chamfer distance and normal consistency.

A.2 Number of eigenvectors

In the top left of Figure 1, we show the number of eigenvectors which captures 99.5% of variance within the ShapeNet dataset. Here, our *eigenSDF* method was trained jointly on all categories. Remaining 3 images present captured variance when our method was trained per category.

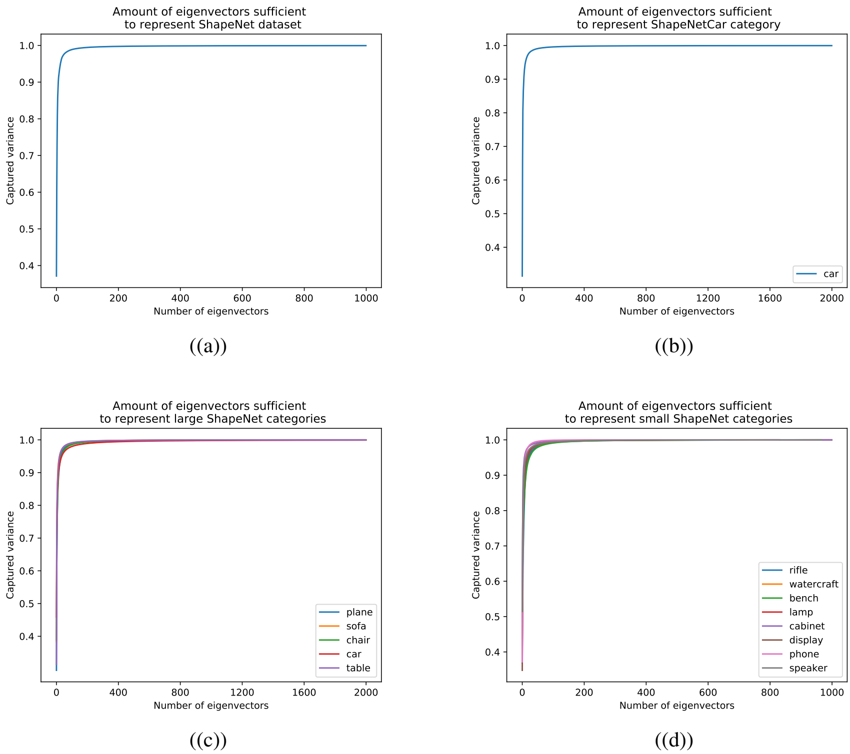

Figure 1: Cumulative fraction of total variance captured by eigenvectors obtained from applying PCA on:
(a) all 13 categories
(b) ShapeNetCars category
(c) large categories, i.e. those having over 3 000 examples
(d) small categories, i.e. those having under 3 000 examples

We further analyze the impact of the eigenvectors on the resolution. Figure 2 shows that the amount of eigenvectors needed decreases with resolution as the variability of shapes in lower resolution is smaller. Small differences further enhance the claim that a more complex benchmark is needed to study the problem of shape inference from a single image.

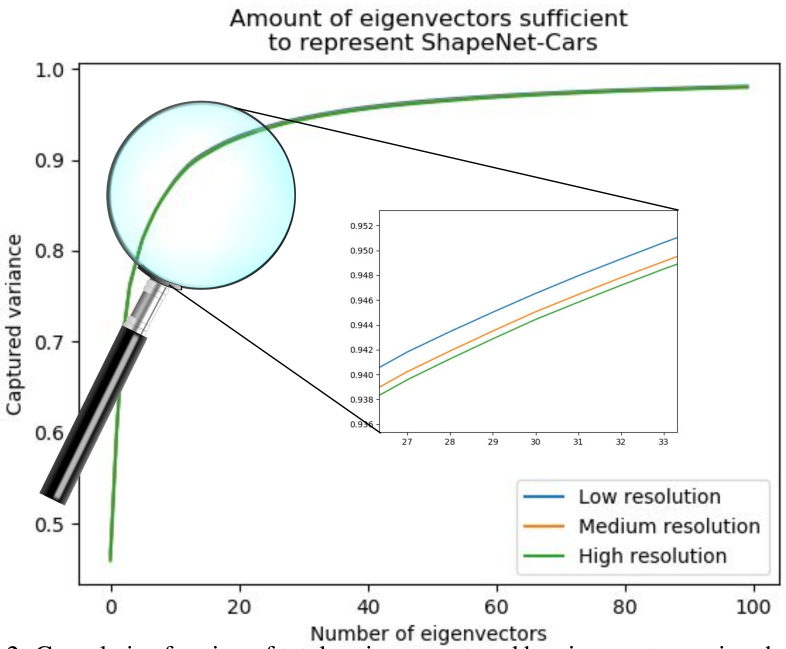

Figure 2: Cumulative fraction of total variance captured by eigenvectors using three different resolutions: low (32^3), medium (64^3) and high (128^3)

A.3 PCA applied to other shape representations

As mentioned in the *Related Work* Section of the main paper, there are 4 commonly used shape representations: voxels, signed distance functions, point clouds and meshes. Because of lack of canonical order or representation it is hard to apply PCA on point clouds or meshes in a straightforward way. Compared to the results of PCA reconstruction using signed distance function as shape representation, applying PCA directly on voxels yeilds poor results, as shown in Figures 3-4.

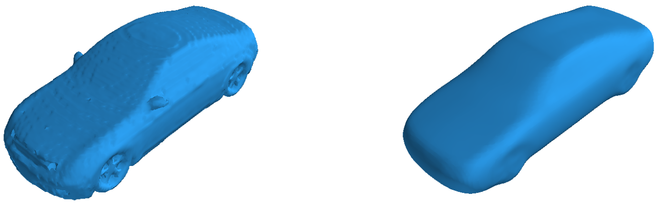

Figure 3: SDF-based PCA reconstruction of a car example from *ShapeNetCars* dataset.
Left: level set of a SDF representing a car with resolution $128 \times 128 \times 128$.
Right: PCA recostruction using 2048 eigenvectors.

Figure 4: Failed attempt of reconstructing a car example from *ShapeNetCars* dataset using voxel-based representation.
Left: voxelized car with resolution $128 \times 128 \times 128$.
Right: PCA recostruction using 2048 eigenvectors.