Appendix

We describe here the color and geometric features used in feature vectors for learning style similarity.

Color and Texture Features

We include the following features to capture color and texture information: average of hue, saturation, and value of five most dominant colors (3), histograms for hue, saturation, and value (32 bins each), and three different resolution Local Binary Pattern bins (10, 14, 18).

To compute the above mentioned features for each model, we use the .obj and .mtl files along with texture images available in folders associated with each object. Specifically, the .obj file gives the information about total number of materials used and which material is assigned to which face. The .mtl file specifies color and texture properties for each material, and the image folder contains texture images used.

We begin by computing the fraction (***F***) of total surface area covered by each material based on the number of faces using it. For this we use the .obj files. Next, for each material specified in the .obj file, we construct two lists using the information in .mtl files. The first list represents color information present in diffuse parameters and second list represents texture images used. Given these two lists, we use the following steps to compute color features:

1. For each material:
   a. If no texture exists, use diffuse parameters to get 32 bins each for the hue, saturation, and value histograms.
   b. If texture exists, use texture image to get 32 bins each for the hue, saturation, and value histograms.
2. Finally, combine the above histograms for all materials to get 32 bins for each hue, saturation, and value histograms.

To compute texture features we use Local Binary Pattern (LBP) on the associated texture images. We use three different resolutions of rotation invariant LBP to construct a 42 dimensions vector. We also include the average HSV of the top five dominant materials based on surface area covered (***F***). These features give us a color feature vector of length 141 dimensions.

Geometric Features

All the geometric features used in our method are based on already existing works. We ask the interested reader to refer to the comprehensive surveys in the literature on the topic of 3D content retrieval and feature based shape similarity retrieval, such as in [Bus05, Iye05]. For completeness, we now briefly describe the geometric features used in our technique.

With the aim to capture both global and local shape properties, we include the histograms for the following geometric features (histogram bins in brackets): shape distribution (128), curvature (gauss, mean, max, min: 128 each), shape diameter (128), light field (470), voxel gradient (192), voxel gradient direction (128), silhouette centroid distances (192), silhouette Fourier descriptor (57), silhouette Zernike moments (108), silhouette D2 descriptor (192), silhouette gradient (192), silhouette gradient direction (96), and shape histogram (192).

Before computing features, all models are oriented in the same direction and scaled to have similar proportions within each object type. Since the input meshes have different resolutions, computing some features (e.g. shape diameter) directly on the 3D mesh results in incomparable or incompatible feature vectors for the learning stage. To rectify this, we use uniformly sampled (with 10,000 samples) surface versions of 3D models to compute the first three features. Specifically, we compute the shape distribution or D2 descriptor as histograms of 128 bins histogram; Gaussian, mean, min and max curvatures histograms with 128 bins each; and shape diameter descriptor with 128 bins. The remaining features above are computed using the volumetric representation of the 3D model by a voxel grid of size 300x300x300. We rasterize the models into binary voxel grids, where a voxel has value 1 if it is on the boundary of the model, and a voxel has a value 0 if it lies elsewhere. To compute voxel gradient from voxel representation, we use 3x3x3 sobel filter along x, y, and z axis. The voxel gradient direction histograms of 64 bins each along x-z and x-y are computed from the voxel gradient. The silhouettes along x, y, and z directions are obtained by projecting the voxel space along the three axes respectively. Silhouettes centroid histogram for each projection (64 bins) is constructed from Euclidian distances between center and boundary points. Additionally, on silhouettes, we compute Fourier, Zernike moments, D2 (between each point on boundary), and silhouette gradient (2D sobel filter) and silhouette gradient direction histograms. Lastly, we compute the shape shell histogram descriptor which is similar to the shape distribution descriptor. The histogram bins for each shell (3) in this case represent the distance of each point to the barycenter of the 3D mesh. Finally, we combine all the features above to give a feature vector of 2587 dimensions.