

Style-Content Separation by Anisotropic Part Scales

Kai Xu^{*†} Honghua Li^{*†} Hao Zhang^{*} Daniel Cohen-Or[‡] Yueshan Xiong[†] Zhi-Quan Cheng[†]

^{*}Simon Fraser University

[†]National University of Defense Technology

[‡]Tel-Aviv University

Abstract

We perform co-analysis of a set of man-made 3D objects to allow the creation of novel instances derived from the set. We analyze the objects at the part level and treat the *anisotropic part scales* as a shape style. The co-analysis then allows *style transfer* to synthesize new objects. The key to co-analysis is part correspondence, where a major challenge is the handling of large style variations and diverse geometric content in the shape set. We propose *style-content separation* as a means to address this challenge. Specifically, we define a *correspondence-free* style signature for style clustering. We show that confining analysis to within a style cluster facilitates tasks such as co-segmentation, content classification, and deformation-driven part correspondence. With part correspondence between each pair of shapes in the set, style transfer can be easily performed. We demonstrate our analysis and synthesis results on several sets of man-made objects with style and content variations.

Keywords: co-analysis, style-content separation, shape analysis, part correspondence, segmentation

1 Introduction

One of the main goals of computer graphics research is to provide effective means for artists and non-experts in geometry modeling to easily create useful digital 3D models. In this paper, we are interested in the creation of novel 3D models of man-made objects which retain the characteristics of a given example set belonging to a certain class; see Figure 1. This problem is particularly challenging when the synthesis process is unsupervised, that is, the given set of objects are not manually tagged with semantic labels and the correspondence among them is unspecified and difficult to compute due to the significant shape variations in the set. Such shape variations are quite typical of many classes of man-made objects.

The set of chairs in Figure 1 (more in Figure 7) clearly contain subclasses such as swivel office chairs, armchairs, folding chairs, and bar stools, as reflected by the geometric “content” of the chairs’ functional parts. At the same time, one can also classify the whole set into clusters based on some notion of “style”, e.g., high vs. low back, long vs. short legs. It seems obvious that the shape parts play a fundamental role in analyzing the set or as we call it in this paper, *co-analysis*. Specifically, the shape parts are the representative entities and not feature points, e.g., there is no obvious feature correspondence between the four legs of a bar stool and a swivel base with five wheels, however they both function as the “leg part”. It is evident that both part identification and making a distinction between the styles and contents associated with the parts represent an essential understanding of the semantics of the set of shapes.

Style-content separation is indeed fundamental to human perception and this ability is reflected in routine tasks such as understanding words (the content) spoken in different accents (the style) and recognizing letters in different fonts [Tanenbaum and Freeman 2000]. An important ability enabled by the separation is the meaningful extrapolation of new data from existing observations, a process of creation which interests us. However, while natural and easy to do by humans, style-content separation is known to be very difficult for machines [Hofstadter 1985]. In this paper, we are interested in developing effective means of style-content separation for



Figure 1: Style-content separation by anisotropic part scales facilitates part correspondence among a diverse set of chairs (top). The separation is shown by the table with rows representing identified styles. The correspondences allow automatic synthesis of novel shapes (shaded in gold) from the example set via style transfer.

co-analysis of a set of 3D models and demonstrate the crucial role such a separation plays during this process and in shape creation.

Works on style-content classification using machine learning techniques have focused on presumed parameterized models [Tanenbaum and Freeman 2000; Wang et al. 2007] or statistical modeling using PCA [Blanz and Vetter 1999; Brand and Hertzmann 2000]. Styles and contents are learned from training data or derived by model fitting and as such, they are generic and data-dependent. In all of these works, a prerequisite is data correspondence and this has been dealt with independently from style-content analysis, e.g., using crude image alignment or designated markers for correspondence. In contrast, we argue that the correspondence problem itself can be the main challenge, particularly for a diverse set of objects such as the one in Figure 1 and others shown in this paper. As a result, style-content separation should be applied at this stage.

Overview. Given a set of 3D mesh models, we co-analyze them at the *part* level. This is motivated by the reduced representative power of feature points for man-made shapes. As the shape style, we consider the individual and relative *anisotropic scales* of the shape parts. Although the anisotropic scale of an individual shape part is an admittedly primitive geometric attribute, when the part scalings in the whole shape are *non-homogeneous*, as exhibited by the chair set in Figure 1, they are already challenging any existing method for shape alignment or correspondence. In our analysis, we use tight oriented bounding boxes (OBBs) of the shape parts as a coarse representation and define a set of OBB graphs to capture the various part compositions of a given shape.

Our style-content separation algorithm proceeds in four steps as illustrated in Figure 2. We elaborate on each step in Section 3.

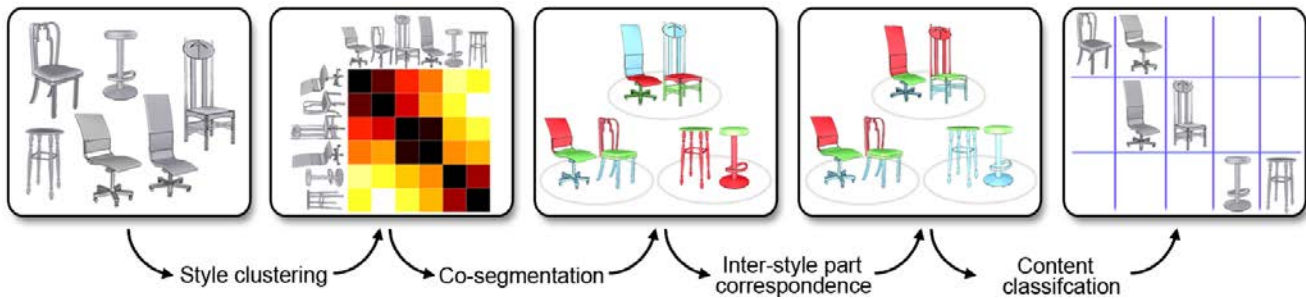


Figure 2: An overview of the co-analysis pipeline.

1. **Style separation:** A style signature called *anisotropic part scale* (APS) can be defined for each part composition of a given shape. Given a part composition, its APS characterizes the individual anisotropic part scales and their differences, which are encoded into the corresponding OBB graph. The signature is given by Laplacian graph spectra, which are permutation-independent, allowing their comparison to be *correspondence-free*. Style separation is achieved by unsupervised learning of the APS signatures via spectral clustering.
2. **Co-segmentation:** This is performed in each resulting style cluster to obtain *intra-style part correspondence*. The approach we take follows the consistent mesh segmentation algorithm of [Golovinskiy and Funkhouser 2009]. Since this method relies on a global rigid alignment, its performance can be greatly improved by style separation which serves to factor out the effect of non-homogeneous part scaling.
3. **Inter-style part correspondence:** This is computed between styles, specifically the co-segmentations derived for the styles, using a novel “deform-to-fit” procedure. The deformation-driven approach allows for partial matching and seeks the best fitting of OBBs through scaling, rotation, and sliding operations. This step is facilitated by co-segmentation since it represents a unified representation for each style, thus reducing the search space for partial matching between styles.
4. **Content classification:** With part correspondence between each pair of shapes in the set, content classification can be performed on corresponding parts using existing shape descriptors, such as the light field descriptor [Chen et al. 2003]. The part-level approach leads to more refined classification results than those obtained by global comparison.

The result of style-content separation is a table where styles are arranged into the rows and contents into the columns. Shape synthesis via style transfer is accomplished by transferring an APS style i (from the i -th style cluster) to any example shape (the source model) in the set, that is, by re-scaling the parts of the source model, through OBB scaling and space deformation, according to style i . The style transfer is rather straightforward as long as there is a part correspondence between the source model and the co-segmentation of the i -th style cluster. Such correspondences are already available after deform-to-fit (step 3 in the pipeline).

Contributions. The main contributions of the paper include:

- A new framework for co-analysis of a 3D set via style-content separation by anisotropic part scales, allowing for effective part correspondence and shape creation via style transfer.
- A novel unsupervised algorithm for style-content separation that is correspondence-free.
- A deformation-driven part correspondence scheme suitable for man-made shapes via “deform-to-fit” of part OBBs.

- Style separation by anisotropic part scales effectively removes non-homogeneous part scaling from the co-analysis equation, facilitating analysis tasks such as co-segmentation, part correspondence, and content classification.

2 Related work

Most works on automatic shape synthesis have focused on organic shapes, such as faces [Blanz and Vetter 1999], the body shapes of humans [Allen et al. 2003; Anguelov et al. 2005] or other creatures [Praun et al. 2001; Kilian et al. 2007]. The typical approach first establishes correspondences between a set of example shapes and then computes variations among the corresponding entities. The common shape representation is a polygon mesh and the corresponding entities are given by the mesh vertices. A statistical model of the variations then allows one to generate a new shape via a blending process. A similar paradigm can be applied to the synthesis of motion [Brand and Hertzmann 2000; Wang et al. 2007; Lau et al. 2009] or images [Drori et al. 2003].

Perhaps the best known part-based shape synthesis is the work on modeling by examples [Funkhouser et al. 2004]. The conceptual design task is controlled by the user while content-driven search of whole shapes or shape parts allows geometry variations to be added via part substitution. Our synthesis approach is different and it seeks to automatically generate “more of the same” from an example set of shapes belonging to a certain class. We specifically learn the part scale styles from the set and synthesize by analogy.

To enable style transfer, the key is to obtain part correspondence through co-analysis. Shape correspondence is a well-studied problem [van Kaick et al. 2010]. Most methods, particularly those for registration or statistical modeling, assume that the shapes to be matched during correspondence do not differ significantly. A few more recent papers [Zhang et al. 2008; Lipman and Funkhouser 2009] attempt to handle large shape deformations, but like most previous works, they focus on organic shapes and rely on extracted surface feature points to define a meaningful correspondence. Recent work of Au et al. [2010] also deals with organic shapes but it matches skeletal features which can be seen to represent shape parts. However, their algorithm, like others, is sensitive to topological changes and non-homogeneous part scalings.

There seems to be a recent trend towards the analysis and processing of man-made objects [Fu et al. 2008; Pauly et al. 2008; Xu et al. 2009]. Of relevance is the work of Gal et al. [2009] on iWires, an interactive editing system for man-made shapes based on shape analysis. Bokeloh et al. [2010] derives procedural models from a single input example via symmetry analysis. The analysis results in a decomposition of the example shape so that new instances can be generated automatically using the procedural rules. Our synthesis is also automatic, however instead of using procedural models

computed from a single input example, we seek to generate “more of the same” shapes from a set of examples via style transfer.

Corresponding man-made shapes is difficult due to the oftentimes large intra-class shape variations including non-homogeneous part scaling and even topological changes. The reduced usefulness of feature points is also an issue, preventing the use of many classical correspondence schemes [van Kaick et al. 2010]. Last but not least, meaningful deformation of man-made shapes, e.g., for editing [Gal et al. 2009; Xu et al. 2009], is challenging since such deformation cannot be free-form [Botsch and Sorkine 2008] and a strong emphasis on shape semantics is necessary.

Much work has been done on mesh segmentation but mostly on individual models [Katz and Tal 2003; Shamir 2008]. Existing works on co-analysis of a set include those on consistent mesh segmentation, which can possibly imply a group correspondence. Shapira et al. [2008] rely on the shape diameter function and Kraevoy et al. [2007] on convexity and compactness to consistently segment organic shapes including humans and animals. Golovinskiy and Funkhouser [2009] resort to global alignment to initialize a clustering scheme to compute consistent mesh segmentation. Global alignment to remove anisotropic scales has been proposed by Kazhdan et al. [2004] to improve shape alignment and retrieval. The key difference between our approach and the last two methods is that our scaling or alignment is *part-aware* and in particular, it is able to successfully handle non-homogeneous part scaling, which is common to man-made shapes [Kraevoy et al. 2008].

Finally, we mention the recent work of Kalogerakis et al. [2010] on learning mesh segmentation and labeling. Assigning semantic labels to a set of consistently segmented shapes does imply a part correspondence. However, they use *supervised* learning, while our co-analysis is entirely *unsupervised*.

3 Style-content separation and style transfer

The input to our co-analysis (Figure 2) is a set of shapes given by polygonal meshes. The mesh models need not be closed manifolds but we assume that each consists of one or more piecewise smooth surface patches. The models used in this paper are man-made shapes from the iWires dataset [Gal et al. 2009] and the Princeton Shape Benchmark [Shilane et al. 2004]; these models all share the geometric characteristics described above. The output of the co-analysis is a table such as the one in Figure 1, the result of style-content separation. In this section, we detail the steps of the co-analysis pipeline. First, we define the OBB graph, which provides a coarse representation of a given part composition of an input shape.

OBB graph. The oriented bounding box (OBB) of a shape part captures the scales of the part along three major directions. These scale attributes define the APS signature used during style clustering. The OBBs also serve as the primitives for the deform-to-fit step which computes inter-style part correspondence. Now we describe how an OBB graph is constructed for a given input mesh.

We first segment the mesh into parts, where many choices of existing mesh segmentation algorithms [Shamir 2008] may be employed. We have chosen the normalized cut approach based on surface concavity, one of the algorithms considered in the random cuts approach to mesh segmentation [Golovinskiy and Funkhouser 2008]. Obviously, determining the appropriate number of segments is not an easy problem. Hence, we simply over-segment each shape and rely on subsequent steps to arrive at the right part count. In our experiments, we over-segment up to $S_{max} = 15$ parts.

We compute the OBB for each part using PCA. The initial OBB graph is then built whose nodes are the OBBs of the mesh parts and

edges encode adjacency between parts. Adjacency is detected simply by performing intersection tests among the mesh parts. During subsequent processing, adjacent parts can be merged resulting in a new OBB for the two parts and a new part composition for the shape. This can be modeled naturally via recursive contraction of the OBB graph. Hence, each derivable part composition through pairwise part merging has an associated OBB graph.

Style clustering. We learn the APS (anisotropic part scale) styles of the input set by clustering and to that end, we need to define an appropriate distance measure between the styles of two shapes. To truly capture how the anisotropic part scales differ between two shapes, a meaningful part correspondence is required. However, our goal is to accomplish this in a correspondence-free manner and allow style-based clustering to facilitate part correspondence.

Clearly, if two shapes have similar anisotropic part scales, then there must be a number n^* such that two size- n^* part compositions (compositions with n^* parts) of the two shapes have the closest match in their corresponding part scales. Thus we need to search through n , the part counts, to find that optimal value. For a particular fixed n , we need to consider all possible part compositions having size n . However, instead of establishing correspondence between the parts, we define a permutation-independent signature so that their comparison is correspondence-free. We design such a signature, which we call the APS signature. The signature is a function of the part composition and we describe it later.

Given two input shapes p and q , for each n within a prescribed range I_{APS} , we search through all part compositions of size n for the two shapes and find the smallest distance between the APS signatures for the part compositions. The *APS style distance* \mathcal{D} between the two shapes is the smallest such distance for all n :

$$\mathcal{D}(p, q) = \min_{n \in I_{APS}, |c_p|=|c_q|=n} \frac{\|APS(c_p) - APS(c_q)\|}{\sqrt[n]{n}},$$

where c_p and c_q are part compositions of shapes p and q , respectively, and $|\cdot|$ denotes the size of a part composition. Note that the distance is weighted to favor larger part count, with α a user-tuned parameter to control the influence. In all of our experiments, we set the range for n to $I_{APS} = [2, 6]$ and α to 1.6.

With the APS style distance computed between each pair of shapes, we perform standard spectral k -means clustering [Ng et al. 2001] and determine the number of clusters k by observing the k -means cost curve and finding the first point at which the curve starts to level off; this is a fairly standard heuristic for determining k [Everitt et al. 2009] and has worked well in our examples.

A last note on style clustering is related to computational efficiency. Instead of exhausting all possible part compositions, we only look through ones that preserve the major symmetries of the original shape [Podolak et al. 2006]; for man-made shapes, this was proven to be an effective pruning heuristic. Under the parameter setting $S_{max} = 15$ and $2 \leq n \leq 6$, there are typically only hundreds of part compositions to consider per shape. Since n is quite small, comparing a pair of APS signatures consists of only solving a few small-scale eigenvalues problems (see below). The time taken to compare two part compositions for $n = 6$ is about 9 milliseconds.

APS style signature. Given an OBB graph representing a part composition of size n , we wish to design a signature that captures the anisotropic part scale information in a way that is independent of part permutations. To this end, we resort to results from spectral graph theory [Chung 1997] on the usefulness of Laplacian eigenvalues in characterizing graphs. The Laplacian graph spectra are permutation-independent and as signatures, they are generally discriminative since it is known that iso-spectral cases are rare.

1D-to-1D			1D-to-2D	
L-shape	T-shape	Cross	End-face	Mid-edge
1D-to-3D	2D-to-2D		2D-to-3D	
End-face	L-shape	T-shape	L-shape	T-shape

Figure 3: Ten possible OBB-to-OBB transformations. 1D, 2D, and 3D OBBs are colored in yellow, red, and blue, respectively.

First, we encode *binary* or *relative* anisotropic part scale information into three graph Laplacians, one per direction of the OBB. Each Laplacian matrix is $n \times n$ and specifically,

$$S_i(u, v) = \begin{cases} -e^{-\|s_i(u) - s_{\theta(i)}(v)\|^2 / \sigma^2} & \text{if } u \sim v, \\ \sum_{z \sim u} -S_i(u, z) & \text{if } u = v, \\ 0 & \text{otherwise,} \end{cases}$$

where \sim is the graph adjacency relation, the Gaussian width σ is set to be the maximal scale difference among all the OBBs, $s_i(u)$ is the scale of the OBB u along the i -th direction ($i = 1, 2, 3$), and this direction is aligned with $\theta(i)$, one of the scale directions of the OBB v ; $s_{\theta(i)}(v)$ is the scale of OBB v along direction $\theta(i)$. Here the alignment of the axes of two adjacent OBBs is computed using the optimal rotation alignment method of [Kazhdan 2007].

The Laplacian graph spectra of S_1 , S_2 , and S_3 contribute $3n$ eigenvalues to the signature. The three spectra are concatenated into a length- $3n$ vector with eigenvalues sorted in ascending order in each spectrum. However, to compare the signatures from two shapes, the three Laplacians need to be brought into proper correspondence. Since the input sets are composed of man-made objects, it is usually possible to compute their upright orientations [Fu et al. 2008]. These upright orientations may be used as a reference to sort the Laplacians. As an alternative, which is what we have chosen to do in this paper, we simply take the ordering which results in the smallest Euclidean distance between the spectral signatures.

The signature constructed so far is seen to vary with change of angles between the OBBs. It is possible to encode intrinsic OBB scales as *unitary characteristics* of the OBBs. To this end, we consider measures of linearity, planarity, and sphericity of an OBB and encode them as vertex weights into graph Laplacians. Our approach is inspired by the work of [Chung and Langlands 1996] which defined certain combinatorial Laplacians with vertex weights. We define three such Laplacians to suit our problem setting, one for each of the above intrinsic scale measures:

$$A_j(u, v) = \begin{cases} -\sqrt{\alpha_j(u)\alpha_j(v)} & \text{if } u \sim v, \\ \sum_{z \sim v} \alpha_j(z) & \text{if } u = v, \\ 0 & \text{otherwise,} \end{cases}$$

where the index j is one of L , P , and S for linearity, planarity, and sphericity, respectively, and $\alpha_j(u) = \text{vob}(u) \cdot c_j(u)$ is the volume-weighted anisotropy of the OBB u , where $\text{vob}(u)$ is the volume of

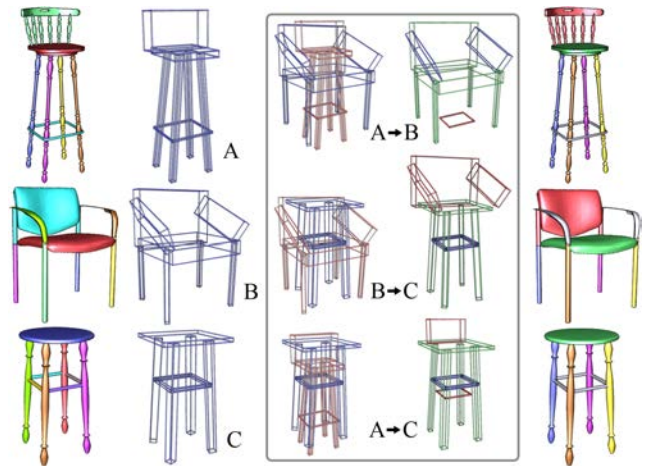


Figure 4: OBB-based part correspondence via deform-to-fit (D2F). Left: three chairs with functional part compositions and their OBB representations. Middle: results of D2F where the red parts are deformed to fit the blue ones; matched parts are colored green; unmatched parts retain their original color. Right: final part correspondences where those parts in the same color are in correspondence while parts in grey are left unmatched.

the OBB u . Assuming that the three scales of u are sorted such that $s_1 \geq s_2 \geq s_3 \geq 0$, then

$$c_L = \frac{s_1 - s_2}{s_1 + s_2 + s_3}, c_P = \frac{2(s_2 - s_3)}{s_1 + s_2 + s_3}, c_S = \frac{3s_3}{s_1 + s_2 + s_3}.$$

We sort the three spectra for $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ in some arbitrary yet agreed-upon order when comparing the signatures of two shapes. The complete APS signature thus contains $6n$ eigenvalues. The distance between two APS signatures is their Euclidean distance.

Co-segmentation. Within a style cluster obtained from the previous style separation step, we perform co-segmentation, which computes a consistent segmentation across all shapes in the style cluster. Such a co-segmentation obviously leads to an intra-cluster part correspondence. Since style separation serves to remove non-homogeneous part scaling from the analysis, co-segmentation under this setting is expected to be an easier problem.

We follow the same approach as [Golovinskiy and Funkhouser 2009] for consistent segmentation. The method first over-segments each shape in the input set (in our experiments, we over-segment each shape to $F_{max} = 25$ segments). Then it performs hierarchical clustering of the resulting segments based on a cost combining two factors: a) *adjacency*, which accounts for how connected two segments from the same shape are for the purpose of mesh segmentation; surface concavity can be used — the more concave the angle between the segments, the larger the cost. b) *correspondence*, which accounts for how well two segments from different shapes correspond; here global rigid alignment is used — the better the two segments can be aligned, the smaller the cost. It is worth noting that it is the rigid alignment step that may render the method unreliable under part stretching, as shown in their work. Style separation provides exactly the right remedy for this very issue.

To adapt the above method to our setting, we need to work at the OBB level, that is, to define the adjacency and correspondence costs based on the OBBs. Since the number of OBBs is typically much smaller than the number of faces, the clustering process is efficient. Given two adjacent segments p_1 and p_2 (initial or clustered) on the




Figure 5: Comparison of co-segmentation results between (a) [Golovinskiy and Funkhouser 2009] and (b) our method. Note the adverse effect of non-homogeneous part scaling on the former and our method’s ability to correct the problem.

same shape, we define the adjacency cost by a convexity measure $(vob(p_1) + vob(p_2))/vob(p)$, where p is the result of clustering segments p_1 and p_2 .

If p_1 and p_2 belong to different shapes, we first find the optimal global alignment between the OBBs of the two *whole* shapes (by aligning their major axes) using rigid ICP [Besl and McKay 1992]. Then we define the correspondence cost as

$$\frac{1}{arob(p_1) + arob(p_2)} \sum_{(f_1, f_2) \in CFO(p_1, p_2)} area(f_1) + area(f_2),$$

where $arob(p)$ is the area of the OBB of segment p and $CFO(p_1, p_2)$ is the set of compatible face pairs between the OBBs of p_1 and p_2 — a pair of faces from these OBBs are compatible if the Euclidean and normal distances between them each falls below a threshold. In our experiments, we set the distance threshold to be 10% of the length of the OBB diagonal for the whole shapes and measure the normal distance by the dot product between normals, setting the threshold at 0.75.

Inter-style part correspondence via deform-to-fit. With co-segmentations computed for the style clusters, we now establish correspondence between these segmentations or part compositions. In general, part correspondence searches for the best match between two given models. We take a deform-to-fit (D2F) approach, finding the best deformation between the matched entities in the source and target models and use certain deformation energy or cost to evaluate the correspondence. Our technique extends the deformation-driven correspondence method of Zhang et al. [2008] to deal with man-made models and to operate at the OBB level. The approach of [Zhang et al. 2008] is that of a pruned priority-driven search. While generally slow due to the exhaustive nature of the approach, it is able to find quality partial correspondences.

The two given models to be brought into correspondence are the co-segmentations from two style clusters and each of them is represented by an OBB graph. We classify each OBB into three types: 1D, 2D and 3D, roughly corresponding to linear, planar, and spherical objects, respectively. To classify the OBBs into these types, we use the three measures c_L , c_P , and c_S defined for APS signatures. Typically, one of the three values is much larger than the rest, uniquely determining the OBB type. However, some OBBs may be regarded as of more than one type, such as the leg part of a stool shown in Figure 1 — since both c_L and c_S are large, the OBB is regarded as both a 1D-type OBB and a 3D-type OBB.

The OBBs and the relation among them define an attributed graph in a straightforward way. Since the number of OBBs is rather small, the search space is small and an exhaustive search is viable [Zhang et al. 2008]. To reduce search time, we select the node with the largest number of edges as the root of the source model. In the




Figure 6: Comparison of content classification on a set of kitchen utensils. Left: with style clustering, D2F, and part-level LFD distances. Right: direct LFD distances measured on whole shapes. Our method does a much better job in separating the spoon, knife, three- and two-pronged fork sets, even under style variations.

target model, we use all the nodes whose types are the same as the root of the source as the candidates for correspondence.

In addition to considering unary characteristics of the OBBs, we also define allowable binary relations; this further reduces the search time during D2F. Specifically, two neighboring OBB nodes have a valid relation if their types can be transformed to fit by a rotation or sliding operation. We define ten possible OBB-to-OBB transformations as shown in Figure 3. These OBB-to-OBB transformations always exist between functional parts of man-made objects and they seem sufficiently representative, which is the very reason why we have identified these transformations. During D2F, before measuring the deformation between two matched OBB sub-graphs, we first validate the OBB-to-OBB correspondence relations; invalid relations lead to pruning of the search subtree.

The quality of each valid correspondence between two sub-graphs is measured by $E_{D2F} = \alpha E_{match} + \beta E_{volume}$, where E_{match} is the volume percentage of unmatched parts and E_{volume} is the sum of squared volume differences of all matched parts; E_{D2F} can be seen as the deformation energy. Throughout our experiments, we simply set $\alpha = \beta = 0.5$. Figure 4 illustrates some results of OBB-based part correspondence through D2F. Results on consistent co-segmentation, which are enabled by style-content separation and inter-style correspondence via D2F, can be found in Figure 8.

Content classification. Most content-based classification algorithms seek to identify the right class for a query [Tangelder and Veltkamp 2008]. For our current problem however, we are faced with an input set already belonging to a class, e.g., chairs, and our goal is to further *refine* the classification. One naturally expects that a global shape descriptor would not work well for this task so that part-level analysis becomes necessary. The part correspondence obtained from the previous step allows us to do just that.

We develop a simple scheme for content clustering. First, we cluster the input shapes based on the part counts in their co-segmentations. Then within each such cluster, we define a distance measure that is the maximum Euclidean distance between the light field descriptors (LFD) [Chen et al. 2003] defined for the corresponding parts. However, before computing the LFDs, we first scale the corresponding OBBs so that they align and the parts within are space-deformed [Cohen-Or 2009]. Results from content classification are reported in Section 4.

Synthesis by style transfer. The most difficult aspect of automatic shape synthesis is the correspondence among the example set. With style separation, co-segmentation, and finally D2F to achieve meaningful part correspondence between each pair of shapes, the



Figure 7: Style-content separation for a chair and stool set (row: style; column: content). Newly synthesized models (in gold) fill the table.

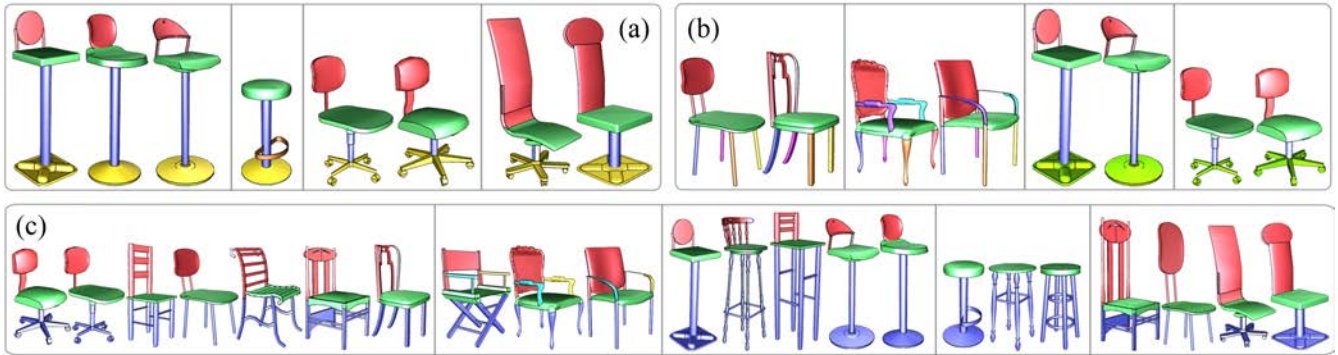


Figure 8: Results of co-segmentation and D2F-based inter-style part correspondence on three subsets (a-c) of the chair and stool set. In each subset, models are grouped by style clusters.

synthesis problem by style transfer is greatly facilitated. With part correspondence, we know precisely how each part in the modified shape needs to be scaled. The only remaining problem is to maintain a proper connectivity between two adjacent parts under possibly different scalings, in particular between the adjoining features.

We have implemented a rather simplistic scheme to handle scaling discrepancies between adjacent parts as it is not the focus of this paper; more sophisticated deformation techniques [Kraevoy et al. 2008; Gal et al. 2009] may be otherwise applied. During style transfer, each part is scaled by deforming its OBB according to the style. For all the synthesis results shown in this paper, we assume that the scalings only cause sliding of the connection curves in an interfacing plane. Then the disconnected connection curves after scaling are simply blended linearly to obtain the new adjoining feature. We obtain the final surface by performing a surface deformation with the blended curve as a handle similar to iWires [Gal et al. 2009].

4 Results

In this section, we show results of style-content separation and shape synthesis. Models used in the example set consist of those from the Princeton Shape Benchmark [Shilane et al. 2004] and Gal’s iWires dataset [Gal et al. 2009], as well as ones obtained via

space deformation to enrich the styles. The sizes of the mesh models range from $5K$ vertices to $40K$ vertices. For an example set of 10 to 20 models with each containing tens of segments (typically ranging from 2 to 25), the typical timing of the various steps of our method are given below. Style clustering: 30-340 seconds; co-segmentation: 15-60 seconds; deform-to-fit part correspondence: 0.5-2 seconds, roughly 0.04 seconds per pair of models; content clustering: 15-40 seconds with most of the time spent on computing LFDs. These times were measured on an Intel(R) Core 2 Duo 2.53GHz CPU machine with 2GB memory.

One of the key points of the paper is that style separation by APS analysis facilitates shape co-analysis and we demonstrate this first. Figure 5 compares our co-segmentation results with [Golovinskiy and Funkhouser 2009], showing that our style separation can properly handle non-homogeneous part scaling. Figure 6 compares content classification results between with and without style clustering, demonstrating the ability of our method to achieve much refined classification. The main reason for such success is that style clustering, co-segmentation, and deform-to-fit, together allow to derive meaningful part correspondence and content analysis at the part level leads to the more refined results.

Figure 7 shows the full set of chairs and stools and the table resulting from style-content separation. Each style cluster occupies




Figure 9: A gallery of style-content tables (row: style; column: content) and synthesis results for six data sets: hammers, airplanes, tables, kitchen utensils, goblets, and humanoids. It is especially interesting to note the ability of our method to successfully content-classify the shapes within a class despite the rather small geometry variations among shapes in that class, e.g., see the hammer and utensil sets.

a row, while content clusters occupy the columns. Newly synthesized models fill the table and in this and all other figures, they are shaded in gold. The table also shows results of content clustering — shapes belonging to the same content clusters are grouped into the same columns. Despite the large style variations in the input example set, our method exhibits success in the results described above. Note that since our D2F procedure is able to obtain partial matching, we can even perform style transfer between two styles which were not brought into full part correspondence. For example, style transfer can be performed between stools and chairs, and

between chairs with and without armrests. A side effect worth noting however is that the newly created models may not respect the styles identified for the original example set. One may also notice that some synthesis results contain distortions, as exemplified by the result given in the second row, eighth column of the table of Figure 7. This may occur due to the use of homogeneous resizing for the matched parts which can possibly lead to distorted segments.

To stress-test our algorithm, we apply it to a diverse set of object classes shown in Figures 9. Note that despite our focus on man-made shapes, we also ran our algorithm on a set of organic (hu-




Figure 10: *Unnatural part correspondence (bottom: A and B with respect to the rest) and synthesis results (top: red background).*

manoid) shapes. While the processing of all other sets were automatic, following the procedures described in the paper, the humanoid set required some manual effort to obtain reasonable results. Six out of the ten input humanoid models were segmented manually since the particular mesh segmentation algorithm we are adopting does not work well on the set. In particular, the segmentation boundaries obtained lack sufficient quality. Manual part stitching was also performed during style synthesis since our simplistic linear blending between the adjoining feature curves works poorly on organic shapes. These two issues indeed point to certain limitations of our current implementation, which we discuss in Section 5. Nevertheless, the results on all the sets in Figure 9 demonstrate success in style-content separation for co-segmentation and content clustering, which represent the core of our contribution.

The vehicle set given in Figure 10 also presents difficulties in terms of initial segmentation and part stitching, for which manual efforts were involved. However, even with such efforts, we encounter a case of unnatural part correspondence, as highlighted by the two models marked *A* and *B*. In particular, the back wheels of these models are not properly matched with any of the wheels of the remaining input models. As a result of the unnatural correspondence, the synthesis results, shown by red background in the table, are also seen as unnatural. The main cause of the problem is that our D2F-based correspondence currently does not alter the part compositions obtained during co-segmentation. Since the latter only works within style clusters, inconsistencies between part compositions across different clusters may occur. Having D2F search through different part compositions can solve the problem but would incur a high computational cost. Alternatively, we could segment the green part of *B* into two parts, a cab and a trailer; this would also fix the problem but the initial segmentation would have been hard to obtain using only geometric cues.

5 Discussion, limitation and future work

In this paper, we deal with the difficult problem of analyzing a set of 3D objects belonging to the same class while exhibiting significant shape variations, particularly in part scales. Existing shape analysis algorithms have been ineffective in dealing with non-homogeneous part scales, and the associated co-analysis problem is particularly challenging under the unsupervised setting. Treating anisotropic part scales as a particular shape style, we propose to apply style-

content separation prior to co-segmentation and correspondence computations, effectively factoring out the part scale issue from the analysis equation. The strength of our approach stems from the power of our style clustering scheme. Being correspondence-free, it succeeds to reveal valuable knowledge about the set, which then paves the way for subsequent analyses including co-segmentation, part correspondence, and refined content classification.

Admittedly, the current work is merely making one step forward in solving the general problem. It is important to note that our co-analysis is not expected to work on an arbitrary input set. The set needs to contain objects in the same semantic class, and with sufficient variety in its anisotropic part scales. In addition, the initial (over) segmentation of the input shapes needs to be sufficiently meaningful. As we can observe from the vehicle and humanoid sets, achieving this automatically is not always easy.

As a means to generate novel shapes, our synthesis method limits itself to creating new variations of an existing example model instead of arriving at new contents by exchanging or shuffling parts of multiple shapes. While the latter may be able to generate more shape varieties, it inevitably has to take care of the part stitching problem. Stitching parts from different models, especially man-made shapes, is quite a difficult problem in general; it may require resolving topological inconsistencies between parts and the general problem is not even well-defined. Our current solution to part stitching is admittedly rather simplistic; tackling the general problem indeed points to an interesting direction for future work.

We are witnessing a rapid accumulation of 3D models yet most of these models are not tagged or segmented and their parts have no semantic labels. We believe that co-analysis of sets of shapes, as we do in this paper, will gain more and more interest, both for indexing and retrieving shapes, and as a tool for generating new variations of existing objects. As shown in this paper, style-content separation can be an effective means to this end. However, anisotropic part scales is only one particular shape style. We believe that our current investigation is only a start to defining and incorporating more shape styles into the shape analysis problem.

Acknowledgments

We first thank the anonymous reviewers for their valuable comments. We are grateful to Aleksey Golovinskiy, Joshua Podolak, Thomas Funkhouser, Michael Kazhdan, and Ding-Yun Chen for sharing their software or source code on consistent segmentation, symmetry detection, as well as spherical harmonics and light field shape descriptors. Thanks also go to Yanzhen Wang and Ariel Shamir for discussions on the topic. This work is supported in part by grants from NSERC (No. 611370), the Israeli Ministry of Science, the Israel Science Foundation, the 863 Program of China (No. 2007AA01Z313 and 2009AA01Z301), NSFC (No. 60773022, 60773020, and 60970094), and the Pre-research funding of National University of Defense Technology (No. JC09-06-01).

References

- ALLEN, B., CURLESS, B., AND POPOVIĆ, Z. 2003. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. on Graphics* 22, 3, 587–594.
- ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., RODGERS, J., AND DAVIS, J. 2005. Scape: shape completion and animation of people. *ACM Trans. on Graphics* 24, 3, 408–416.
- AU, O. K.-C., COHEN-OR, D., TAI, C.-L., FU, H., AND ZHENG, Y. 2010. Electors voting for fast automatic shape correspon-

- dence. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 29, 2, 645–654.
- BESL, P. J., AND MCKAY, N. D. 1992. A method for registration of 3-d shapes. *IEEE PAMI* 14, 2, 239–256.
- BLANZ, V., AND VETTER, T. 1999. A morphable model for the synthesis of 3D faces. In *Proc. SIGGRAPH*, 187–194.
- BOKELOH, M., WAND, M., AND SEIDEL, H.-P. 2010. A connection between partial symmetry and inverse procedural modeling. *ACM Trans. on Graphics* 29, 4, 104:1–10.
- BOTSCH, M., AND SORKINE, O. 2008. On linear variational surface deformation methods. *IEEE Trans. Vis. & Comp. Graphics* 14, 1, 213–230.
- BRAND, M., AND HERTZMANN, A. 2000. Style machines. In *Proc. SIGGRAPH*, 183–192.
- CHEN, D.-Y., TIAN, X.-P., SHEN, Y.-T., AND OUHYOUNG, M. 2003. On visual similarity based 3D model retrieval. *Computer Graphics Forum (Proc. EUROGRAPHICS)* 22, 3, 223–232.
- CHUNG, F. R. K., AND LANGLANDS, R. P. 1996. A combinatorial laplacian with vertex weights. *J. Comb. Theory Ser. A* 75, 2, 316–327.
- CHUNG, F. R. K. 1997. *Spectral Graph Theory*. AMS.
- COHEN-OR, D. 2009. Space deformations, surface deformations and the opportunities in-between. *Journal of Computer Science and Technology* 24, 1, 2–5.
- DRORI, I., COHEN-OR, D., AND YESHURUN, H. 2003. Example-based style synthesis. In *Proc. IEEE Conf. on CVPR*, 143–150.
- EVERITT, B. S., LANDAU, S., AND LEESE, M. 2009. *Cluster analysis*. Wiley.
- FU, H., COHEN-OR, D., DROR, G., AND SHEFFER, A. 2008. Upright orientation of man-made objects. *ACM Trans. on Graphics* 27, 3, 42:1–7.
- FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. on Graphics* 23, 3, 652–663.
- GAL, R., SORKINE, O., MITRA, N. J., AND COHEN-OR, D. 2009. iWIRES: an analyze-and-edit approach to shape manipulation. *ACM Trans. on Graphics* 28, 3, 33:1–10.
- GOLOVINSKIY, A., AND FUNKHOUSER, T. 2008. Randomized cuts for 3D mesh analysis. *ACM Trans. on Graphics* 27, 5, 145:1–12.
- GOLOVINSKIY, A., AND FUNKHOUSER, T. 2009. Consistent segmentation of 3D models. *Computers & Graphics (Proc. of SMI)* 33, 3, 262–269.
- HOFSTADTER, D. 1985. *Metamagical themas*. Basic Books.
- KALOGERAKIS, E., HERTZMANN, A., AND SINGH, K. 2010. Learning 3D mesh segmentation and labeling. *ACM Trans. on Graphics* 29, 3, 102:1–12.
- KATZ, S., AND TAL, A. 2003. Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Trans. on Graphics* 22, 3, 954–961.
- KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2004. Shape matching and anisotropy. *ACM Trans. on Graphics* 23, 3, 623–629.
- KAZHDAN, M. 2007. An approximate and efficient method for optimal rotation alignment of 3D models. *IEEE PAMI* 29, 7, 1221–1229.
- KILIAN, M., MITRA, N. J., AND POTTMANN, H. 2007. Geometric modeling in shape space. *ACM Trans. on Graphics* 26, 3, 64:1–8.
- KRAEVOY, V., JULIUS, D., AND SHEFFER, A. 2007. Model composition from interchangeable components. In *Proc. of Pacific Conference on Computer Graphics and Applications*, 129–138.
- KRAEVOY, V., SHEFFER, A., SHAMIR, A., AND COHEN-OR, D. 2008. Non-homogeneous resizing of complex models. *ACM Trans. on Graphics* 27, 5, 111:1–9.
- LAU, M., BAR-JOSEPH, Z., AND KUFFNER, J. 2009. Modeling spatial and temporal variation in motion data. *ACM Trans. on Graphics* 28, 5, 171:1–10.
- LIPMAN, Y., AND FUNKHOUSER, T. 2009. Möbius voting for surface correspondence. *ACM Trans. on Graphics* 28, 3, 72:1–12.
- NG, A. Y., JORDAN, M. I., AND WEISS, Y. 2001. On spectral clustering: analysis and an algorithm. In *NIPS* 14, 849–856.
- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. 2008. Discovering structural regularity in 3D geometry. *ACM Trans. on Graphics* 27, 3, 43:1–11.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Trans. on Graphics* 25, 3, 549–559.
- PRAUN, E., SWELDENS, W., AND SCHRÖDER, P. 2001. Consistent mesh parameterizations. In *Proc. SIGGRAPH*, 179–184.
- SHAMIR, A. 2008. A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6, 1539–1556.
- SHAPIRA, L., SHAMIR, A., AND COHEN-OR, D. 2008. Consistent mesh partitioning and skeletonization using the shape diameter function. *The Visual Computer* 24, 4, 249–259.
- SHILANE, P., MIN, P., KAZHDAN, M., AND FUNKHOUSER, T. 2004. The princeton shape benchmark. In *Proc. IEEE Conf. on Shape Modeling and Applications*, 167–178.
- TANENBAUM, J. B., AND FREEMAN, W. T. 2000. Separating style and content with bilinear models. *Neural Computation* 12, 6, 1247–1283.
- TANGELDER, J. W. H., AND VELTKAMP, R. C. 2008. A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications* 39, 3, 441–471.
- VAN KAICK, O., ZHANG, H., HAMARNEH, G., AND COHEN-OR, D. 2010. A survey on shape correspondence. In *Proc. of Eurographics State-of-the-art Report*.
- WANG, J. M., FLEET, D. J., AND HERTZMANN, A. 2007. Multifactor gaussian process models for style-content separation. In *Proc. of Int. Conf. on Machine learning (ICML)*, 975–982.
- XU, W., WANG, J., YIN, K., ZHOU, K., VAN DE PANNE, M., CHEN, F., AND GUO, B. 2009. Joint-aware manipulation of deformable models. *ACM Trans. on Graphics* 28, 3, 35:1–9.
- ZHANG, H., SHEFFER, A., COHEN-OR, D., ZHOU, Q., VAN KAICK, O., AND TAGLIASACCHI, A. 2008. Deformation-driven shape correspondence. *Computer Graphics Forum (Proc. SGP)* 27, 5, 1431–1439.