



Figure 2: A family of sofas of varying dimensions and functions reuses parts and their relations (IKEA sofa collection EKTORP). Although the models have very different geometric representations, at a high-level, they have similar style. We focus on extracting structure from low-level geometric descriptions to characterize such a notion of style leading towards a semantic understanding of shapes.

able.

In classical shape modeling, the user remains in charge of manually recognizing and asserting any higher-level structural constraints, and resolving any dependencies. To reduce the burden on designers and artists, we need methods that automatically recognize structural properties and invariants of a shape or shape families, that understand their interdependencies, and that assist the user in creating structurally plausible shapes efficiently. Thus, the processing pipeline is necessarily structure-aware throughout.

In this survey, we review methods that support structure-aware shape processing. We focus in particular on techniques designed for *man-made* objects. These objects often have simple structural invariants across larger classes of shapes, such as symmetry, coplanarity, orthogonality, or regular arrangements, allowing us to characterize whole families of shapes concisely. Even though the low-level geometric descriptions of each individual object might be detailed and complex, the subparts of the objects share many such local and non-local relations across a larger family of related shapes. Thus, they can be compactly represented and parameterized as few part geometries along with the relationship among the parts, e.g., as a graph or hyper-graph.

A historical perspective. In the context of computer aided geometry, the use of algebraic methods to handle constraints was pioneered in the 1980-90s. We refer the readers to the survey by Hoffmann et al. [Hoffmann and Joan-Arinyo 2005] for key aspects of geometric constraint solving with characterization of the different methods used for constraint driven modeling. Instead, in this work we focus on the challenges arising from having to deal with real unorganized models, which naturally require geometric data understanding (e.g., 3D scans; non-manifold and heterogeneous quality 3D models).

Up until the mid-2000's, geometry modeling in graphics was dominated by techniques focusing on low-level shape analysis and processing. Typical examples include smoothing, subdivision, surface reconstruction, and Laplacian -based detail-preserving editing. A notable exception is the work of Funkhouser et al. [Funkhouser et al. 2004] on modeling by example, where new shapes are generated by part composition based on a database of existing shapes (see Figure 3). This modeling paradigm still plays the dominant role in novel shape creation today, with the latest variants including structure-aware part shuffling [Kraevoy et al. 2007], probabilistic synthesis [Chaudhuri et al. 2011; Kalogerakis et al. 2012], set evolution [Xu et al. 2012a], etc.

The analyze-and-edit approach introduced in the iWires work [Gal et al. 2009] has set the tone for numerous works on structure-



Figure 3: Modeling-by-example proposed a novel model synthesis possibility by directly composing parts retrieved from a database of 3D shapes [Funkhouser et al. 2004]. As 3D model repositories continue to grow, such a design possibility has increasingly gained research focus.

preserving editing, e.g., [Zheng et al. 2011; Lin et al. 2011; Bokeloh et al. 2012]. The analysis phase often takes the center stage as it presents the taller challenge of shape or structural understanding. Such works include inverse procedural modeling [Stáva et al. 2010; Bokeloh et al. 2010] and symmetry-driven generative analysis [Pauly et al. 2008; Wang et al. 2011]. Note early work by Stiny [Stiny and Gips 1971] on shape grammars has been an inspiration in this area, with current works focusing on recovering such grammars directly from the input examples.

Symmetry indeed plays a key role in structure-aware processing of man-made objects, since symmetries are abundant in these objects and from a functional point of view, symmetric structures are expected to share the same functionality. Methods for structural symmetry detection [Mitra et al. 2006; Podolak et al. 2006; Simari et al. 2006] have been made accessible to the geometric modeling community and since then, a large body of work has been developed on symmetry-driven shape processing (cf., survey [Mitra et al. 2012]). An attempt to understand functionality of man-made objects was the notion of upright orientation [Fu et al. 2008], which can be seen as one of the early efforts in structural analysis of man-made objects. Both symmetry and upright orientation represent early attempts at extracting high-level semantic information from geometry.

Part analysis lies at the core of structure-aware shape processing. Extracting parts from a shape is the classical segmentation problem (c.f., survey [Shamir 2008]). Equally important is the part correspondence problem, not only between a pair of shapes, but among a shape collection (c.f., survey [van Kaick et al. 2011b]). Early works on both the problems focused exclusively on local geometric analysis including those on geodesics and curvature. An influential effort was made by the ShapeAnnotator framework [Atene et al. 2009] which helps the user create semantic, part-based shape ontology. Recent developments have taken the data-driven approach, in particular by learning structural invariants from a set of exemplars. Such a *co-analysis* approach started with the work of Golovinskiy and Funkhouser [Golovinskiy and Funkhouser 2009] on consistent segmentation and gained momentum on several fronts including style content separation [Xu et al. 2010], joint segmentation [Huang et al. 2011], spectral clustering [Sidi et al. 2011; Kim et al. 2012a], active learning [Wang et al. 2012], co-




Figure 4: An example of a structured shape (a) that is decomposed into parts. Each part controls a portion of geometry (b), which we call part geometry. The parts (c) have parameters and a constraint energy that controls the parameters as well as the decomposition itself. The example shows a pairwise relations (a part graph). In this particular case, it has a Markovian structure (only neighbors interact). Such graphs would for example reflect the structure of enforcing connectivity or smoothness across neighboring pieces.

abstraction [Yumer and Kara 2012], and discovery of functional substructures [Zheng et al. 2013]. Going beyond object modeling, Fisher et al. [Fisher et al. 2011; Fisher et al. 2012] analyze scene data using spatial contexts and apply the results for novel scene synthesis.

Organization. We organize this survey on structure-aware shape processing techniques by the underlying structure models as primary sorting criterion. We believe that formalizing our understanding of “structure” in shapes is the intellectual key challenge in creating more powerful and generally applicable modeling tools. Beside discussing the broad types of structure analysis approaches, we also summarize key application areas including smart acquisition, structure-aware model editing, novel model synthesis, and design space exploration. We conclude by listing some of the key challenges to be solved and their potential implications.

2 Overview of Structure

Starting with early scientific efforts to relate object form to its intended functions, the inter-relation between form and function has played an important role in shape understanding and design. In nature, the geometric forms, such as the skeleton, of many organisms can be explained by physical considerations arising from their environment, operational efficiency, and functional utility. For example, the horns and shells of animals are often spiral-shaped due to growth pattern and rate considerations [Thompson 1992]. Hence, it is expected that functionally related natural objects would share a similar form, or as we shall call in this survey, structure.

In man-made artifacts, such similarities can be even more pronounced since most man-made objects are designed to serve singular functions. Structures of man-made objects arise due to physical and economic, i.e., cost, constraints during design and manufacturing, or more pragmatically, due to semantic considerations. But, what is a *structure*?

The Oxford dictionary defines structure as

(Definition): The arrangement of and relations between the parts or elements of something complex.

In this view, structure constitutes a collection of parts and how the parts are mutually related. It should not be confused with the notion of structure in structural engineering. In this survey, we follow this abstract notion for composite objects. Such a characterization is not surprising given that most man-made objects (e.g., chairs, tables, lamps, shelves) are a constellation of parts, where the constellations are often characteristic of the corresponding model collection. We assume that the global geometry is captured by a composition of (abstract) *parts*, each of which has *parameters* that define the parts, and a set of *relations* among these parts that impose structure on the composite. A unique feature of the setup is that the relations

can arise among parts restricted to any particular model, or more generally, across different but functionally related models (e.g., a collection of chairs).

2.1 Modeling Structures

We describe a generic model for structuring shapes, which we later use to compare and contrast apparently different structure-aware shape processing techniques, and better understand their scope and design choices. A shape \mathcal{S} can be seen as a collection of parts, their parameters, and most importantly, the relations that characterize the arrangement of the parts (see Figure 4).

Parts and parameters. A part P_i of a shape is a logical entity of semantic significance that controls the appearance of part geometry. Note that in this abstraction, parts are not necessarily disjoint, i.e., they can overlap. Further, each part has a finite set of parameters that affect the shape of the part. Note that unlike in traditional geometry processing, by part we do not necessarily mean a surface patch arising from segmentation. Instead, a part can simply be an abstraction for a region of the object and act as a proxy for a semantic part. Figure 4 shows a schematic example of a shape that is decomposed into parts with multiple parameters per part.

Generally speaking, a part a vector of all the vertices forming a shape space (e.g., as in [Yang et al. 2011]), a feature curve (e.g., as in [Bokeloh et al. 2009; Gal et al. 2009]), a bounding box (e.g., as in [Xu et al. 2010; Ovsjanikov et al. 2011]), a fitted proxy such as a box or generalized cylinder (e.g., as in [Zheng et al. 2011; Xu et al. 2011; Xu et al. 2012a]), a surface or volumetric segment [Shamir 2008], or a variational proxy, etc. Note that in most cases the choice of parts automatically determines the choice of the respective parameters. Further, in some cases, the parts can be completely specified by the user, i.e., semantic parts can be provided as part of the input as an input template. Most methods covered in this survey operate on parts which are meaningful components of a shape [Shamir 2008], e.g., a leg of a chair, a table top, a wing of an airplane, or a window over a facade, etc.

Relations. Relations capture how parts, and hence their parameters, are correlated. Such relations can be between a pair of parts (i.e., a pairwise relation) or among a set of parts (i.e., higher order relations). The relations are the key element behind any structure. Formally, relations can be represented mathematically by a *constraint energy* E that must be zero for a valid structure (relations enforced as hard constraints) or that should be minimized (soft constraints). Let us consider few examples:

- In a constraint-based modeling setup, the relations would require elementary, pairwise relations such as parallelity, coplanarity, bilateral symmetry, etc.

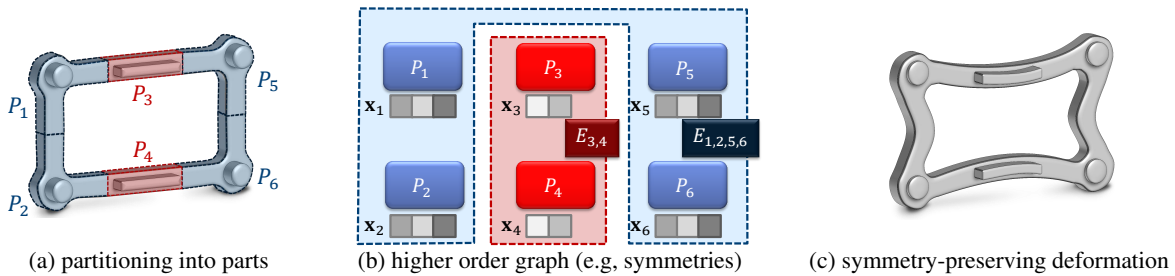


Figure 5: We can also use higher-order relations that couple larger cliques of parts (a), including non-local (non-Markovian) relations. In this example, we show a part (hyper-) graph (b) that aims at preserving symmetry. It couples symmetric pieces using hyper-edges (red/blue), leading to symmetric parameter changes. Such a setup can for example be used for symmetry preserving shape deformation (c).

- Relations can also link a set of parts by considering higher-order relations. For example, the windows on a building facade can be arranged on a 2-parameter regular grid. In general, symmetry relations with more than two elements involved will lead to higher-order relations. Figure 5 shows an example – the blue parts are coupled by a 4-ary symmetry energy.
- A set of relations can capture the part configuration of a functional model collection, e.g., a set of chairs. Such constraints should capture both coupled part-level continuous variations and also discrete changes like chairs with or without arms. Thus, the relations capture not only part-arrangements, but also how they can vary in a coupled fashion, i.e., a corresponding deformation model.

2.2 Discovering Structures

For most methods covered, the input shapes come in the form of low-level geometric representations, such as polygonal meshes, thus obscuring the underlying structure of the shapes. For example, a 3D scanner returns a collection of points, possibly with normals; or, models from 3D model repositories (e.g., Trimble 3D warehouse, Turbosquid) often come as polygon meshes (manifold, or polygon soups), mostly without access to any semantic information. A key challenge is to discover the structure from a single model, or a set of related example shapes (i.e., a model collection). Thus, we have to determine (i) the parts, (ii) their parameters (i.e., representation), and most importantly, (iii) their relations. These are in fact mutually dependent, so ideally they should be learned in conjunction to each other.

A key aspect of structure-aware processing is to *discover* or learn the underlying structure from input examples — it is an analysis problem we refer to as structural analysis (Section 4). Starting from generic assumptions of how parts and relations are formed, the methods semi- or fully-automatically extract a structure model from one or a set of shapes. In Section 4.2, we categorize the main strategies used to discover structure models, which include symmetry-driven analysis, statistical modeling, and *co-analysis*, which learns a structure model from a shape collection.

2.3 Structure-aware Shape Processing

Shape or part structures, once extracted, can be used for a range of structure-aware shape processing tasks (Section 5). A recurring theme in these tasks is structure preservation, e.g., during shape editing and synthesis. The working hypothesis is that preserving structures also ensures plausibility of the manipulated or synthesized models. Complementary to local differential operators, structure models might also provide *non-local relations* and thus couples distant parts of the shapes. Interestingly, focusing on struc-

ture over low-level details allows us to couple models even with large geometric and topological variations. Thus, structure or part constellations become the key focus, with geometry playing only a secondary role. Preserving structures in course of scanning, editing, modeling, or even design exploration, then leads to meaningful and intuitive shape processing. The structures provide high-level handles for model manipulation that often reflect object semantics, thus vastly simplifying the modeling interface.

3 Types of Structures

In order to discover a structure, one has to solve three problems: (i) identify the relevant parts which is akin to a segmentation problem; (ii) identify the relevant parameters or representations for the parts which is akin to part parameterization problem; and (iii) identify relations (i.e., correlation) among the parts both inside a shape and also across different shapes in a model collection. The challenge is perform these steps in conjunction.

We now categorize different approaches mainly based on how each of parts, their parameters, and their relations are discovered (see Figure 6). In each case, depending on the richness of the input data, broadly there are three options: (i) Models are manually prescribed by the user; (ii) models are extracted based on some prescribed model amounting to model-fitting; and (iii) models are learned directly from the data in an unsupervised or supervised setting.

3.1 Identifying Parts

The first task is to devise a scheme for splitting objects into parts. Ideally, such parts should relate to semantic object parts, rather than being constructed from low-level geometric details.

User specified parts. The simplest idea is to let the user explicitly create or specify parts in the input data sets. For example, in a classic modeling system, the user creates parts such as spheres, cubes, etc. To this end, many approaches use information such as connected components of a triangle mesh [Jain et al. 2012; Zheng et al. 2013] or the scene-graph hierarchy provided in a database of 3D models [Fisher et al. 2011].

Parts are obtained according to a fixed model. In this case, a fixed segmentation model is constructed and applied. Popular choices are for example the detection of crease lines [Gal et al. 2009], of objects with a certain symmetry properties [Mitra et al. 2010; Mitra et al. 2013; Bokeloh et al. 2010; Kalojanov et al. 2012], or generic shape segmentation approaches [Shamir 2008].

Parts are discovered directly from the data. Instead of using a model that is completely based on a priori assumptions, it is also possible to define a meta-model that describes how the part structure should be extracted using a learning approach. The difference

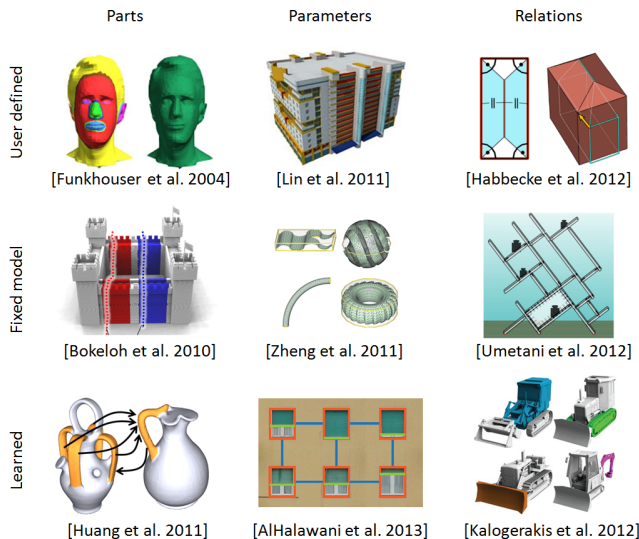


Figure 6: Examples of various methods that define or discover parts, parameters, and relations, i.e., extract structure for the respective collections.

here is that the concrete extraction model will be adapted according to some set of training data. If this process requires user annotations, this is a supervised learning technique [Anguelov et al. 2005; van Kaick et al. 2011a; Kalogerakis et al. 2010; Sunkel et al. 2013]. Unsupervised methods use clustering in order to find structural similarities in the input data, typically a collection of shapes belonging to the same family, automatically [Golovinskiy and Funkhouser 2009; Xu et al. 2010; Sidi et al. 2011; Huang et al. 2011; Hu et al. 2012; Kim et al. 2012a].

3.2 Identifying Parameters

Having established a part decomposition, the next task is to extract its parameters.

User specified parameters. This option is frequently used in procedural modeling schemes. By permitting the user to write scripts or shape grammars, custom parameters can be specified interactively. There are also more traditional systems (from the 1980s) where the whole modeling process is driven by specifying a scene hierarchy with user defined parameters used for communication between modules that instantiate geometry [Stiny and Gips 1971].

Parameters are obtained according to a fixed model. There are a large number of methods where the set of parameters is fixed a priori. For examples, parts are arranged using rigid transforms [Funkhouser et al. 2004; Gal et al. 2009], isometric variations [Huang et al. 2009], or restricted based on fabrication requirements [Yang et al. 2011]. Very frequently, parts can be arranged with rigid motions or general homogeneous transformation matrices. Controlling parts with axis-aligned scaling of the bounding boxes is another option [Zheng et al. 2011]. Parametric models also include specifying the geometric properties of fixed primitives (radius of a sphere, width/height/depth of a cuboid).

Parameters are learned directly from the data. In the most general setting, the set of parameters is not known a priori but the degrees of freedom have to be extracted from training geometry. This is more complicated and structurally very related to the extraction of global constraints from training data, which we discuss more in detail in the next subsection. Common models for learning parameter sets from data are dimensionality reduction techniques, such as

principle component analysis [Duda et al. 2000], or its non-linear variants (kernel PCA, spectral embedding techniques with custom kernels) [Shawe-Taylor and Cristianini 2004].

3.3 Identifying Relations

The key aspect of structure is relation among parts. However, to relate parts often amounts to establishing a global correspondence among the parts. This is challenging as models often have significant geometric and topological variations even with the same model collection. In such scenarios, point-to-point correspondence is ill-defined. An alternative is to assign a probabilistic or fuzzy correspondence among shape collections [Kim et al. 2012a]. Alternatively, explicit correspondence extraction can be avoided as proposed by Ovsjanikov et al. [Ovsjanikov et al. 2011] who analyze model shape collections as a manifold of characteristic descriptors. The descriptors are chosen such that they smoothly depend on the shape. Assuming continuous variability, exploring the shape space can be mapped to navigating a non-linear manifold.

Relations can capture how a part relates to itself (i.e., symmetry) and thus can be modeled as an unary term. More interestingly, pairwise part relations can be captured as binary relations; while more generally interaction between set of n parts can be captured as n -ary relations. Since n -ary relations can be difficult to track computationally, they are often split into lower-order relations, which is strongly related to graphical models in machine learning (c.f., [Koller and Friedman 2009]). Pairwise relations lead to graphs where nodes represent parts and their parameters and edges correspond to the constraint energy induced by pairwise relations. Higher order models correspond to hypergraphs, where each hyper-edge might connect more than two parts. If the graph is sparse and only connects locally neighboring parts, we have a Markov-random field (assuming a soft constraint energy that has a probabilistic interpretation, for example as the log-likelihood of an actual probability density). For the interpretation as a (classic) probabilistic graphical model, we have to assume that the set of employed parts is fixed a priori. The graphical model in Figure 4 is an example for a Markovian model with pairwise edges. In contrast, Figure 5 is not a Markovian model, because there are non-local couplings. Furthermore, and independently of this fact, higher-order edges are used (the graph is a hypergraph).

User specified relations. In the context of computer-aided design and sketching interfaces, a number of systems have been developed where the user can specify general geometric constraints that are checked and maintained automatically by the modeling system [Zeleznik et al. 2007; Igarashi et al. 1999]. This is not the central topic of our survey; we will focus on automatic methods.

Relations are obtained by a priori model. In this case, the constraints are fixed, independent of the data. In many cases, these are physical constraints: For example, properties such as connectivity (dynamic stability), center of mass (for balance), or constraints on the maximum stress within the material (for static stability) can be imposed [Umetani et al. 2012; Whiting et al. 2012]. Deformation models also fall into this category — they prescribe priors on differential properties of a deformation function acting upon a piece of geometry. These soft-constraints are given a priori by analogy to physical processes such as the elastic or plastic deformation of rods, shells, or solids.

Relations are discovered from the data. The final, most interesting, case is to learn constraints from data. This can be done with machine-learning approaches of different strength: A weakly learned model would employ a rather rigid a priori assumption and only apply it to example data in order to obtain concrete rules. An example of such an approach is the learning of shape grammars

based on local similarity [Mitra and Pauly 2008; Bokeloh et al. 2010]: Here, the invariant is fixed a priori (all shapes must be similar to an example shape within local regions). The learning is restricted to computing the rules that arise from applying this model to a concrete piece of example geometry. A much more flexible model is for example the famous morphable face model of Blanz et al. [Blanz and Vetter 1999]: It only assumes low-dimensional linear subspaces for models and learns the geometric variability automatically from corresponding 3D scans. Such a model can also be extracted directly using abstracted part geometry and proxies [Ovsjanikov et al. 2011; Kalogerakis et al. 2012].

4 Structure Analysis

In this section, we discuss recent techniques developed for learning structure models, which is an analysis task. The learned structures are then utilized for a variety of processing tasks, in particular, shape synthesis. Works on structure-aware processing are surveyed in Section 5.

4.1 Challenges in Structure Discovery

Although, in the previous sections, we independently considered parts, parameters, and their mutual relations, ideally they should be learned simultaneously. This makes the problem challenging, especially in the fully unsupervised setting. For example, what are good object parts depends on their mutual relations, but the mutual relations can only be explicitly extracted when the parts along with their correspondences are known. In order to break this cyclic dependency, either training data is used, or the user prescribes candidate parts, or auxiliary model information (e.g., keywords, tags, scene-graphs) are usually used. Further, in many settings, parts can have significant variations across even related objects, and hence a corresponding deformation model has to be additionally extracted or factored out.

Instead of focusing on the variations (i.e., deformation), the general strategy is to look for consistency and patterns. This is motivated by statistical learning theory, which requires the repeated observation of objects originating from a common process. Only then is it possible to gather statistics about the outcome of the process and predict future events of this type. It is, in general, impossible to learn information from one process by observing a second, completely unrelated process [Duda et al. 2000]. In terms of the part-based model defined above, any function of a subset of parts can be such a property if it (i) is observed repeatedly, and (ii) shows a significant statistical relation across these observations.

Thus, the definition of structural invariants requires identifying objects that correspond to the same underlying process. Please note that this abstract notion of correspondence does not necessarily refer to correspondence of points or regions in shapes. Correspondence is meant as an abstract association of properties that are related. For example, this could include generic properties such as elastic properties of local surface elements, or non-local properties such as the overall surface area of an object. Broadly, there are three main approaches, which we describe next.

(i) *Global point-wise correspondences*: These techniques consider a family of shapes with global, point-wise correspondences between the considered shapes. In our setting, the objects have only one part, with a number of corresponding parameters essentially representing a shape space (e.g., [Kilian et al. 2007]). If applied directly to geometry (for example, every vertex of a mesh being a parameter), this implies that the topology of the created or edited shapes cannot be changed from what was determined in the analysis phase. A classic example is a morphable model (providing a lin-




Figure 7: Models, even from the same family, can have large variations in geometry and topology. As a result, the notion of point-to-point correspondence is not always well defined in this context. Hence, Kim et al. [Kim et al. 2012a] introduce the notion of fuzzy correspondence, which can then be used for interactive exploration of heterogeneous model collections.

ear subspace of possible configurations) (e.g., [Allen et al. 2003]). Physically based models such as elastic deformation also use global correspondences, but the constraint energy is not learned but rather specified a priori (e.g., [Umetani et al. 2012]).

(ii) *Partial point-wise correspondences*: These techniques are more general and do not require a one-to-one identification of all points in all shapes, but permit association across partial data. Hence, they correspond to techniques that permit multiple parts, where each part has its own parameter set. Only parts of compatible types, however, can be put into correspondences. Typical examples of part-based models are texture synthesis, (inverse) procedural models, or regularity-based shape resizing. Further, in the notion of model collections, Kim et al. [Kim et al. 2012a] introduce the notion of fuzzy correspondence, where instead of a point-to-point correspondence a fuzzy or probabilistic notion of correspondence is proposed (see Figure 7).

(iii) *Correspondence across derived properties or “descriptors”*: These models do not identify actual surface points, but derived properties that summarize larger portions of the shape. Such models are typically *discriminative* [Kumar and Hebert 2003]. They can detect changes of a property of a shape in a collection of shapes, but they are typically not *generative*; they cannot generate new shapes from scratch, just from knowing the derived property, because there is a loss of information in computing this property. Typical examples are methods that identify objects by descriptors, possibly combined with discriminative, supervised learning algorithms that structure the space of descriptors according to labeled examples. Such learning is for example performed in the data-driven shape segmentation approach by Kalogerakis et al. [Kalogerakis et al. 2010].

4.2 Discovering Structures

We now survey works developed for structural analysis. There are different ways to classify these techniques. In terms of the learning mechanism, they can be supervised, unsupervised, or semi-supervised. In terms of what input a technique takes, they can be based on analyzing a single shape or studying a whole collection, i.e., co-analysis. From a computational point of view, we can examine the techniques by what optimization strategies are employed.

In what follows however, we categorize the surveyed methods in terms of the broad strategies employed in the analysis; a key ques-

tion is what prior knowledge is relied upon. Supervised techniques rely on training data [Kalogerakis et al. 2010] or user-specified templates. We focus on unsupervised approaches.

4.2.1 Symmetry as Source of Structure

Symmetry is a purely geometric notion, i.e., it can be computed based purely on geometric analysis without any additional prior knowledge. Yet it carries powerful semantic information and plays a critical role in structural analysis of man-made objects, whether the input is an individual shape or a shape collection. In the case of analyzing individual models, symmetry provides the key source of structure. The most important reason is that symmetric parts in an object are expected, in fact, they were designed, to perform the same function. Hence, detected symmetries provide the first cues for functional analysis. Symmetry-driven analysis is primarily based on the prior knowledge that symmetric parts are to be grouped together [Wertheimer 1923].

Symmetry is the absence of information (i.e., redundancy) within a shape. Typically, a group \mathcal{G} of allowed transformations are specified. Most often, \mathcal{G} is a subset of rigid, similarity, or isometric transforms. In order to have a symmetry, the group \mathcal{G} of operations must not change at least a subset $\mathcal{P} \subseteq \mathcal{S}$ of a geometric object \mathcal{S} . In this case, we have found a (partial) symmetry: All of the operations $g \in \mathcal{G}$ are symmetry transformations that keep \mathcal{P} invariant.

In general there are two different types of symmetries: The weak form are *instances* or *building blocks* [Kalojanov et al. 2012]: The set of operations \mathcal{G} refers to all possibilities of exchanging parts of the object with one another (for example by rigid transformations) without affecting the shape itself (in this case, \mathcal{G} is not a subgroup of the rigid motions). The stronger form are *regular transformations* [Liu et al. 2004; Pauly et al. 2008]: In this case, the transformations acting on the geometry themselves form an algebraic group. This means, that not just the geometry is redundant in the sense of transformations of the same piece showing up repeatedly, but the way these building blocks are instantiated are regular themselves. Slippability [Gelfand and Guibas 2004] can be considered as a special case (continuous case) of such regularity with respect to rigid motions.

Various methods have been proposed for detecting symmetry [Alt et al. 1988; Gal and Cohen-Or 2006; Loy and Eklundh 2006; Martinet et al. 2006; Mitra et al. 2006; Podolak et al. 2006; Simari et al. 2006; Ovsjanikov et al. 2008; Pauly et al. 2008; Bokeloh et al. 2009; Xu et al. 2009a; Raviv et al. 2010; Xu et al. 2012b]. For a detailed survey and in-depth discussion of symmetry, see for example the recent surveys for images [Liu et al. 2010] and for 3D geometry [Mitra et al. 2012].

Symmetry is attractive as a structure model because it is agnostic of the actual geometry that shows a certain symmetry; it can therefore abstract from the concrete shape and can be used to formulate complex, high-level structure assumptions on larger classes of geometry. For example, symmetry has been used to complete partial shapes by inferring the most symmetric one that is consistent with the observation [Thrun and Wegbreit 2005]. Mitra et al. [Mitra et al. 2007] have proposed a method to fully automatically make shapes more symmetric.

Symmetry can also be used as an invariant for controlling object deformation. The iWires system is based on the assumption that a free-form edit of an object should maintain the original symmetry properties of the input shape [Gal et al. 2009] (Figure 14). The interaction and structure detection is based on salient line features [Kent et al. 1996; Ohtake et al. 2004], which are named ‘wires’ [Singh and Eugene 1998]. Similar thoughts have been used




Figure 8: Symmetry hierarchy provides a structural organization of an object’s parts. Parts associated with a node are colored blue. Each node represents either a grouping by symmetry (green) or assembly (red) by connectivity.

by Bokeloh et al. [Bokeloh et al. 2011; Bokeloh et al. 2012] and Xu et al. [Xu et al. 2011], among others, in order to deform objects in a (symmetry) structure preserving manner.

Hierarchical nesting of symmetries has also been explored for shape representation and modeling [Martinet 2007; Wang et al. 2011]. In particular, Wang et al. [Wang et al. 2011] constructs a symmetry hierarchy from a given man-made object (Figure 8). The object is assumed to have been segmented into meaningful parts. Inter-part symmetries as well as self-symmetries of the parts are first detected. Symmetry and (part) connectivity relations are encoded into a graph representation. Then recursive graph contraction is performed to form a hierarchical representation. The order of the contraction operations is determined by a set of hand-crafted rules which reflect the perceptual laws of grouping and the compactness principle (i.e., striving for the simplest representation). The resulting symmetry hierarchies are shown to reflect object semantics and support structure-aware hierarchical shape editing. However, consistency between the hierarchical structural representations obtained on a set of related shapes is not ensured.

4.2.2 Statistical Models of Structure

In the image domain, assembling images from local neighborhoods has been used in non-parametric texture synthesis. The underlying idea is to statistically model structure locally by assuming a Markov-random field (MRF). Instead of modeling global statistics over complete images only image pieces are considered and their potential connection to neighboring pieces. Non-parametric texture synthesis was introduced by [Efros and Leung 1999] and inspired further research for images [Wei and Levoy 2000; Hertzmann et al. 2001; Kwatra et al. 2003; Kwatra et al. 2005; Simakov et al. 2008; Wei et al. 2008]. Also, the idea has been tried for 3D geometry: [Bhat et al. 2004; Sharf et al. 2004; Zelinka and Garland 2004; Lagae et al. 2005; Lai et al. 2005; Nguyen et al. 2005; Zelinka and Garland 2006; Zhou et al. 2006; Merrell 2007; Merrell and Manocha 2008; Chen and Meng 2009; Merrell and Manocha 2009; Alhashim et al. 2012]. MRF-based approaches are powerful tools for modeling local texture but lack the ability to produce globally consistent results (e.g., closing loops). This is due to the MRF assumption that captures only local statistics while the global model structure is ignored. Assembling globally consistent pieces would result in a NP-hard optimization problem. Typically, non-parametric methods damp this problem by minimizing inconsistencies rather than avoiding them completely.




Figure 9: Unsupervised co-segmentation of a highly varied set of container objects [Sidi et al. 2011].

Kalogerakis and co-workers learn a generative probabilistic model describing a component-based structure from a pre-segmented set of shapes [Kalogerakis et al. 2012]. Their key observation is that structural variability heavily depends on the particular style of a shape and its components. Sailing ships, for example, typically differ vastly from cargo ships and components such as sails or containers naturally point to a particular style. This relationship is modeled with a probabilistic model consisting of observable random variables and latent variables that model the cause for these observations. The observable variables can be extracted from the input shapes such as the number of components, shape descriptors, and adjacency information (part parameters). The style of a shape and the style of sub components are modeled as latent variables. The model is organized hierarchically: The shape style conditions on the number of components and their style which in turn conditions on the observable variables describing a component. Additionally, lateral conditional dependencies are learned between observed random variables that capture relationships between attributes of different components.

A different approach is to model the global structure directly and thereby restrict the shape space to globally consistent shapes. In traditional procedural modeling the structure is encoded into a procedure that the user has to provide [Stiny and Gips 1971; Parish and Müller 2001; Wonka et al. 2003; Müller et al. 2006; Lipp et al. 2008]. Various methods were proposed to guide procedural modeling to achieve the anticipated result [Talton et al. 2011; ?]. Inverse procedural modeling poses the problem of creating procedures from examples. In an early work [Hart and Flynn 1997] find fractal representations for input shapes. Recently, this line of research was picked up again for vector graphics [Ijiri et al. 2006; Ijiri et al. 2008; Huang et al. 2009; Yeh and Měch 2009; Stáva et al. 2010] and shapes [Aliaga et al. 2007; Mitra and Pauly 2008; Bokeloh et al. 2010; Bokeloh et al. 2011; Bokeloh et al. 2012; Vanegas et al. 2012; Talton et al. 2012].

4.2.3 Co-analysis of Model Collections

In recent years there have been increasing interests in the co-analysis of sets of shapes. The premise is that understanding a shape, and in particular its parts, can only be reliable by observing a set of shapes possessing the same semantic or functionality, and not from a single observation. While this claim is straightforward, that more information can be extracted by simultaneously analyzing a set, it is unclear how to exploit the distributed knowledge of a set efficiently to achieve a coherent and reliable analysis of structures.

Unsupervised co-analysis only relies on the weak prior that the in-

put set of shapes are related, that is, they belong to the same object class, e.g., chairs, airplanes, or lamps. The central problem of co-analysis is that of *co-segmentation*, where the challenge is to simultaneously segment all the shapes in the input set in a consistent manner (Figure 9). Thus, besides partitioning the shapes into segments, we also obtain a labeling of the segments across the set, where the parts with the same label serve the same semantic purpose, albeit possibly being geometrically dissimilar.

Existing attempts to co-segmentation can be broadly classified into supervised and unsupervised methods. In the supervised setting [Kalogerakis et al. 2010; van Kaick et al. 2011b], a training set with enough pre-analyzed shapes is assumed to be given. The training set is then used to probabilistically label a set of unknown shapes. Although supervised methods are not strictly speaking a co-analysis (since the shapes are not simultaneously analyzed), the result of the labeling leads to a consistent segmentation for shapes in each object category.

The unsupervised setting is more challenging, since the entire knowledge must be extracted from the input set itself [Golovinskiy and Funkhouser 2009; Xu et al. 2010; Sidi et al. 2011; ?; Hu et al. 2012]. It should be noted that no method, be it supervised or unsupervised, can guarantee a perfect co-segmentation of a set (see Figure 9), since the geometry alone cannot always fully convey shape semantics. In particular, no descriptors can capture all possible geometric variations of a part.

Early work by Golovinskiy and Funkhouser [Golovinskiy and Funkhouser 2009] rigidly pre-aligns all the shapes in the set and then clusters the shape faces according to an underlying graph. The graph links faces that are adjacent in the models and faces that are close-by after the alignment. The resulting clusters provide a natural co-segmentation of the shapes.

Xu et al. [Xu et al. 2010] factor out the scale variation in the shape parts by first clustering the shapes into different styles, where style is defined by the anisotropic part scales of the shapes. In this manner, they are able to co-segment shapes with more variability, compared to [Golovinskiy and Funkhouser 2009], and also synthesize new shapes by transferring (the part proportion) styles, as shown in Figure 10. However, the technique is designed to only deal with one specific style.

Sidi et al. [Sidi et al. 2011] pose the co-segmentation problem as that of clustering segments resulting from an initial over-segmentation of the set of shapes using diffusion maps. Compared to the two previous works [Golovinskiy and Funkhouser 2009; Xu et al. 2010], the clustering is performed in a space of shape descriptors rather than on the spatial coordinates of the shapes themselves. This allows the handling of corresponding parts differing in pose, location, and even cardinality, factors which would challenge any technique based on spatial alignment or direct clustering of shape geometry. In addition, the descriptor clustering approach allows to exploit a key enabling feature of the input set, namely, third-party connections. Even if two shapes possess parts that are significantly dissimilar, we can still establish a link between them if there are other parts in the set (third parties) that create such a connection, resulting in a successful co-segmentation. In a follow-up work, Hu et al. [Hu et al. 2012] improve the feature modeling aspect of the algorithm. Instead of concatenating the different features into one feature descriptor, they pose and solve a subspace clustering problem in multiple feature spaces.

Huang et al. [Huang et al. 2011] develop a joint segmentation method where a set is used to assist in the segmentation of *individual* shapes. They can handle rich shape variations by using shape descriptors, but without aiming at a consistent segmentation



Figure 10: Style-content separation [Xu et al. 2010] 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.

of the entire set. Thus, strictly speaking, it is not a co-segmentation method.

Semi-supervised methods can be viewed as unsupervised methods assisted by input coming out of the set, typically though user interaction. In the recent work by Wang et al. on active learning for co-segmentation [Wang et al. 2012], the system automatically suggests constraints which would most effectively refine the results while the user interactively adds constraints as appropriate. Specifically, the user adds either a cannot-link constraint or a must-link constraint between currently obtained segments. The constraints are propagated to the set and the co-segmentation is refined. Their work indeed demonstrates that relatively few user constraints can lead to almost perfect co-segmentation results. Their work also made available the COSEG dataset for evaluating algorithms on co-segmentation.

4.3 Structure from Physical Constraints

Some recent techniques analyze an input man-made object based on domain-specific priors arising from physical or manufacturing considerations. For example, physical laws (e.g., gravity, stability) can be used as algebraic constraints or relations to guide geometric optimization towards form-finding [Umetani et al. 2012; Whiting et al. 2012]. Similarly, manufacturing constraints can determine structural preferences. For example, in the context of architectural buildings with glass facades [Yang et al. 2011], planar faces (e.g., quads) are more desirable as they are cheaper to build. A corresponding constraint amounts to planarity constraint for the respective faces of the surface mesh (e.g., for quad meshes). Other constraints can be horizontal floor lines, or equal height for different floors, etc. Finally, motion considerations for example in automata for toys, the respective parts (e.g., gears, links, etc.) should satisfy certain contact and kinematic constraints, again specified in the form of algebraic constraints [Mitra et al. 2010; Zhu et al. 2012].

5 Structure-aware Shape Processing

Sematic shape processing is a long-standing goal in model creation and synthesis. Structure, either prescribed as a part of the input or discovered by analyzing model collections, gets us closer to this goal. We defined structure as relations between parts (and their parameters) with the relations encoded as constraints or energy functions. As a result deformation, synthesis, and design exploration can all be cast as various forms of constrained-editing.

Optimization Strategies. Various optimization strategies have been employed for structural-aware processing as instances of constrained optimization: (i) greedy progressive projection to the multiple constraints [Gal et al. 2009]; (ii) gradient descent [Umetani et al. 2012; Whiting et al. 2012]; (iii) tangent space navigation using a level-set approach [Yang et al. 2011; Bokeloh et al. 2012]; (iv) modal analysis by studying second order information of the constraints [Yang et al. 2011]; or sparse basis (e.g., using basis pursuit) to locally represent the solution space [Habbecke and Kobbelt 2012; Deng et al. 2013]. The main considerations in deciding among these various options are simplicity of implementation; speed and interactive-response; the order of the constraints; and the dimension of the shape space determined by the structure model.

5.1 Structure-aware Shape Editing

In this section, we discuss a few example systems that aim at *editing* existing shapes. All of these have in common that they base their structure model on an analysis of a single input shape. The discovered structure is then maintained as a soft or hard constraint. This could also be viewed as removing degrees of freedom in comparison to general, unconstrained shape modification, with the goal of making it easier to perform plausible changes.

We consider three classes of examples: (i) As a baseline, we consider traditional free-form deformation. (ii) We then discuss methods that extend these ideas towards the preservation of more general relations, such as salient shape features or symmetry. We pick the “iWires” system [Gal et al. 2009] as a canonical example of such approaches. (iii) The last example concerns methods that are actually able to change the topology of the object. We discuss algebraic regularity models as one example of such methods [Bokeloh et al. 2012].

5.1.1 Baseline: Free-form deformation

If we assume a triangle mesh as input, the simplest method for editing is certainly performing low-level mesh operations (moving vertices, possibly creating and deleting triangles). Obviously, this provides the most degrees of freedom and the least structural constraints. In practice, editing of shapes in this way is impossible for any but the simplest 3D models.

Instead of moving individual vertices, several methods have been proposed to perform higher-level edits that affect multiple vertices at the same time in a useful way. The first methods were free-form deformations [Sederberg and Parry 1986; Coquillart 1990] that use a low-dimensional, band-limited, volumetric basis to impose smooth, low-frequency deformations to the geometry (such as a B-spline / Bezier basis of $3 \times 3 \times 3$ control points). The vertices $\mathbf{x} \in \mathbb{R}^3$ are deformed by a function $f: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that is composed out of such low-frequency basis functions $b_i: \mathbb{R}^3 \rightarrow \mathbb{R}^3$:

$$f(\mathbf{x}) = \sum_{i=1}^n \mathbf{x}_i b_i(\mathbf{x})$$

The underlying structure model is that high-frequency details of the geometry should be preserved while the low-frequency bands are




Figure 11: Left: Setting up a differential soft-constraint energy, every one-ring neighborhood is a part P_i , and every vertex is a parameter x_i , shared by multiple parts. From the configuration of vertices in the one-ring neighborhood, an estimate for differential surface properties is obtained. Right: Deformation result for an elastic deformation model.




Figure 12: Non-homogeneous resizing protects salient regions [Kraevoy et al. 2008].

subject to direct editing. According to our taxonomy, the parts in this approach are the scalar basis functions b_i that represent localized, low-frequency degrees of freedom in the global deformation field f . The parameters are the coefficients x_i (which are 3D vectors). The part-geometry is the set of triangles that have vertices within the support of the part b_i . There is no constraint energy — the user has to adjust all parameters manually. Structure is only implicitly imposed by using only a few, low-frequency basis functions.

The requirement to adjust all control points of a low-frequency deformation basis explicitly creates some avoidable burden to the designer. More recent free-form deformation techniques use physics-based priors: The method uses an elastic energy as constraint energy. Additional handles are added to give the user control. In order to formulate elastic deformation models, differential properties of the deformed shape are computed and deformations that stretch or bend the surface are penalized using a soft constraint energy. If the input is a triangle mesh, the differential properties are typically read off a one-ring neighborhood in the mesh [Sorkine and Alexa 2007], so that the parts of such a model are all one ring neighborhoods and the parameters are the vertex positions of the mesh. Figure 11 illustrates this notion (left), along with a result from an elastic free-form deformation of the Stanford bunny model (right).

There are a large number of variants of this idea, dating back to the seminal paper of Terzopoulos et al. [Terzopoulos et al. 1987]. A survey is provided by Botsch and Sorkine [Botsch and Sorkine 2008]. From a birds-eye-view, the structure model always involves a constraint energy that penalizes deviations from the original shape. By formulating the penalty in a differential domain, local details are preserved more strongly than low-frequency shape properties.

5.1.2 Structure-Aware Deformation

Free-form deformations have a *local* and *non-adaptive* way of preserving structure: The shape of local pieces is preserved independently of the content (lack of adaptivity). Further, no global rela-

tions are considered; only a tendency towards low-frequency bending arises implicitly from chaining differential parts. Both of these aspects open up room for new methods.

Local adaptivity. Local but adaptive deformation is considered for example by [Kraevoy et al. 2008], see Figure 12. Their approach uses a differential free-form deformation energy that prefers axis-aligned stretch. Further, it estimates the “vulnerability” of local regions by looking at differential properties (curvature, slippage analysis [Gelfand and Guibas 2004]), and the elasticity of the model is adaptively reduced in vulnerable regions. Xu et al. [Xu et al. 2009b] introduce adaptivity by adjusting the deformation penalties locally to match the slippability properties of the object (again based on slippage analysis); see Figure 13. This creates a deformation behavior that often mimics the behavior of mechanical systems (joints, cylinders etc.).




Figure 13: Possible joint locations are extracted by a local slippability analysis and subsequently used in enabling joint-aware deformation [Xu et al. 2009b]. Note that structure in the form of the kinematic motion chains is discovered by analyzing the input model. In both of these cases, the same concept of parts and parameters is used. The only difference is how the constraint energy is created. In both cases, the constraints only react to local differential properties of the deformation field, but the behavior is more adaptive than in a standard deformation approach.

Non-local relations. Global relations are used by iWires [Gal et al. 2009], see Figure 14. The method first detects crease lines in a triangle mesh, which the authors call “wires”. These elements are the parts of the deformation model. The parameters are the vertices of the wires. Among these parts, all salient properties of Euclidean geometry (parallelity, orthogonality, different types of symmetry) are enumerated in the analysis phase. These properties are then considered invariants of the deformation. Again, the user can control the deformation by handles, and at first a conventional elastic solution is computed. However, after that, the system tries to restore the invariants iteratively by a greedy algorithm that bends the wires such that they still meet the user constraints but better preserve the original structure. Parallel lines should remain parallel, and objects that were symmetric under a Euclidean (rigid) transformation should maintain this property.

A similar idea is examined by Zheng et al. [Zheng et al. 2011], see Figure 15. Here, the parts are object-aligned bounding boxes of shape components (obtained from segmentation). Again, Euclidean invariants (symmetries) are used to propagate edits to affect all symmetric elements similarly.

Habbecke and Kobbelt [Habbecke and Kobbelt 2012] observe that fully constrained systems are of limited interest in the context of interactive modeling as they lack any design freedom. Hence they focus on under-constrained systems. They linearize the constraint functions (i.e., energies) and examine their nullspace (see also [Bokeloh et al. 2012]) as an instance of the Cardinality Minimization Problem and efficiently solve the optimization using ideas from compressed sensing (see Figure 16). Unlike image-based con-




Figure 14: The *iWires* system models shape structure by symmetry relations within the object (specifically, among line features named “wires”). These properties are preserved during interaction [Gal et al. 2009].



Figure 15: Controllers (shown in cyan), either a cuboid or a generalized cylinder [Zheng et al. 2011], serve as high-level deformation handles for structure-aware deformation of man-made shapes [Xu et al. 2011]. The numbers indicate the number of controllers needed, showing that a small number of controllers usually suffice for the modeling tasks.

strained reconstruction approaches, their proposed system interactively supports manipulation even with hundreds of vertices and constraints.

5.1.3 Topological Changes

All of the methods so far apply homeomorphisms to the original shape, i.e., continuous, bijective mappings that preserve the topology of the input. More recently, structure-aware editing methods have appeared that are able to change the topology of the input.

The modeling by example system [Funkhouser et al. 2004] introduced the concept of model synthesis as putting together parts from different models to form a novel and coherent shape (see Figure 3). The challenge, however, is how to define a family of shapes that is similar to the input but actually changes in composition. A line of work by Bokeloh et al. [Bokeloh et al. 2011; Bokeloh et al. 2012] uses the invariant of maintaining the regularity properties of the input shape. Regular patterns are a special kind of symmetry property: A part of the shape is repeated multiple times in a regular fashion, for example by a rigid transformation. Regularly repeating elements are important structural features of a shape; they can com-




Figure 16: Modeling and editing a roof structure in a constrained-optimization setup. The vertices are relaxed in the analysis stage, and later automatically updated by the editing system [Habbecke and Kobbelt 2012].




Figure 17: Regularity preserving shape editing by algebraically enforcing translative regularity patterns. Left: Degrees of freedom. Right: Shape variants [Bokeloh et al. 2012].

municate direction, grouping, emphasis, layout, order, etc. Other important structural and visual properties of a shape such as straight lines, planes, cylinders, and spheres can be represented as continuous regular patterns. Using this regularity structure as invariant for shape deformation the resulting shape will resemble the original structure but with different proportions.

Both methods focus on discrete and continuous translational patterns but differ fundamentally in their representation. In [Bokeloh et al. 2011] an elastic deformation field is augmented with anisotropic elasticity constraints penalizing deviation orthogonal to translational patterns while allowing stretch in pattern direction. Non-pattern regions are considered to be salient features and are assigned a high stretch resistance. Thereby, occurring stretch will be preferably diffused within regular patterns. Additionally, the method extracts elements from discrete 1-parameter groups and re-structures the object by inserting or removing elements in order to minimize the distortion in pattern direction for discrete patterns. Similar anisotropic energies have been proposed in [Kraevoy et al. 2008] allowing uniform scaling but minimize non-uniform scaling in salient regions or in [Xu et al. 2009b] using local slippage to simulate mechanical joints.

The model of [Bokeloh et al. 2012] follows a stricter interpretation of maintaining regular patterns (see Figure 17). Instead of constraining an elastic deformation field in a least-squares sense to preserve regularity the method parameterizes shape deformation directly over regular patterns. Each pattern is represented by a small number of variables allowing only elementary changes such as moving the entire pattern or changing the number of elements. The method extracts all possible degrees of freedom by analyzing the nullspace of a linear system of hard constraints linking adjacent patterns and thereby guaranteeing a consistent shape. The resulting model typically offers a small number of changeable parameters that correspond to global changes in the shape and represent the algebraic structure w.r.t. regularity. Quadratic programming is used to combine the convex linear space of shapes with least-squares user constraints and regularization. The key difference to [Bokeloh et al. 2011] is the inability to distort patterns at all. Users can only select valid shapes from the shape space which can be a great advantage in some applications but also be a disadvantage in others.

Formally, in [Bokeloh et al. 2012], a part in their model is a region of geometry that has a fixed regularity pattern, i.e., (i) the same geometry shows up repeatedly, (ii) only differing by a rigid transformation, and (iii) these transformations form a commutative group (or a large enough excerpt of a commutative Euclidean symmetry group of rigid motions). By demanding maximal groups and area, the part decomposition is uniquely defined. Commutative symme-

try groups are isomorphic to integer (or continuous) lattices; therefore, they can be characterized by simple repetition counts, which are the parameters of the model. The constraint energy then just aims at minimizing stretch to regularize the solution.

A different way of structure-aware deformation with topological changes was proposed by Lin et al. [Lin et al. 2011]. They explicitly deal with resizing of irregular architecture where the regularity-based approaches would not work. The input mesh is decomposed into a hierarchy of parts (bounding boxes) manually defined by the user. Part parameters are also defined by the user and specify if a part must stay intact, can be replicated/deleted, or can be scaled. The method analyzes a compatibility graph between adjacent parts and constructs a series of 1D sequences which can be easily resized in a structure-aware manner and constrain the remaining sequences.

5.2 Structure-aware Shape Synthesis

In this section, we survey methods for structure-aware modeling that aim at creating new variants of shapes based on structural similarity to one input shape or a collection of shapes. Methods have been proposed that learn relations of part decompositions from a single exemplar using a fixed model to define parts like regular patterns [?; Mitra and Pauly 2008; ?; ?] and/or partial symmetry [?; Bokeloh et al. 2010]. Other methods operate on segmented shape collections to learn part parameters and relations between them for a class of shapes [Xu et al. 2010; ?; ?].

5.2.1 Variation from Single Example

Inverse procedural modeling. Bokeloh et al. [Bokeloh et al. 2010] introduced inverse procedural modeling, wherein parameters of a procedural model is extracted starting from a single example (see Figure 18). Algorithmically, a shape grammar is constructed from a single input shape that produces r -similar shapes meaning that every r -neighborhood of geometry in the output maps exactly to an r -neighborhood in the input shape. The algorithm decomposes an input shape into building blocks at symmetric regions that allow docking of different building blocks at these cut lines forming so called *docking sites*. These docking sites can be systematically found by investigating partial symmetries of the input shape. A docking site follows the boundary of an r -symmetric region, thereby having a corresponding docking site due to symmetry properties, and separates the model into two disjoint pieces. Intuitively, the non-symmetric part can be replaced by the corresponding non-symmetric part enclosed by the corresponding docking site. The method samples symmetry transformations, computes all docking sites, and combines them into a context-free shape grammar. Although the context-free grammar is actually a limitation, it does enable fast synthesis of similar models without requiring to recompute all available docking sites in each step. Subsequently, Kalojanov et al. [Kalojanov et al. 2012] investigated the space of r -similar




Figure 18: Inverse procedural modeling: Shape grammar computed from input shape (red) synthesizes variations of input (grey)[Bokeloh et al. 2010].



Figure 19: Structure-preserving retargetting of irregular 3D architectures [Lin et al. 2011]. The original input model is shown at the bottom left corner.

shapes further. Intuitively, one can overlay all docking sites and fracture the model simultaneously into micro-tiles. As it turns out, this is a canonical decomposition that can construct the entire space of r -similar models. Both methods are examples for fixed part parameters since every part remains rigid.

Earlier, Merrell [Merrell 2007] used a fixed grid to decompose a model into cubic parameterless parts, analyzes geometric compatibility between these parts along the cut boundaries, and synthesizes new variants of the input shape by recombining pieces with compatible boundaries, similarly to texture synthesis. The method was further extended to handle more complex grid structures as in for example the Sierpinski Tetrahedron[Merrell and Manocha 2008].

Structure retargetting. In the context of facade modeling, the work of Lin et al. [Lin et al. 2011] mentioned above performs retargetting of irregular 3D architectural structures including those of facades. Bao et al. [Bao et al. 2012] generate multiple variations starting from a single facade design. An input facade design is generated by hierarchically segmenting and labeling an input image. The user can then manually specify constraints that should be maintained in any variation, e.g., alignment constraints between windows. Subsequently, facade variations are generated for different facade sizes, where multiple variations can be produced for a certain size reminiscent of traditional media re-targetting approaches. Computing such facade variations is based on interleaving heuristic search with quadratic programming. Yeh and co-workers propose a method to learn neighborhood relationships of tiles in a regular grid [Yeh et al. 2013] that allows the use of multi-tiles (larger regions spanning over multiple grid cells) and takes statistical properties into account that were observed in the example data (see Figure 20).

Photo-inspired synthesis. Xu et al. [Xu et al. 2011] introduce an algorithm for the modeling of 3D man-made objects where the user draws creative inspiration from a single photograph. Model creation is supported by utilizing an available set of 3D candidate models. Specifically, the user creates a digital 3D model as a geometric variation from a 3D candidate; see Figure 21. The variation is obtained by automatically deforming a retrieved 3D candidate to fit the photographed target under the guidance of silhouette correspondence. The set of candidate models have been pre-analyzed to possess useful high-level structural information, which is preserved by the geometric variation so that the final product is coherent with



Figure 20: Tile synthesis by utilizing a probabilistic graphical model learned from examples [Yeh et al. 2013].



Figure 21: A matrix of photo-inspired creations [Xu et al. 2011] (rows: 3D model candidates; columns: photographic inspirations). Structures of the 3D candidates are preserved.



Figure 22: Evolution paths leading to the generation of a chair (bottom) with a new structure [Xu et al. 2012a]. Note that new topologies can be generated by the crossover operator. The ancestors of the final chair belong to sets of shapes that are evolved together, i.e., a set evolution.

its inherited structural information readily usable for subsequent model refinement or processing. On the downside, the modeling technique is unable to create new part structures. Or, structures based on abstracted geometry can be used to directly enable intuitive image manipulation of indoor scenes [Zheng et al. 2012].

5.2.2 Synthesis from Shape Collection

With style-content separation achieved via co-analysis of a set of related shapes, Xu et al. [Xu et al. 2010] perform straightforward style transfer by altering the scales of the shape parts, filling in a style-content table with the transfer results. New variations are generated which possess varying part proportion styles, while keeping the part structures entirely intact; no new part structures are created.

Ovsjanikov et al. [Ovsjanikov et al. 2011] implicitly extract such a morphable model by searching for a low degree of deformation to explain an input model collection using a template shape (as collection of bounding boxes) to specify a representation. Later, Kim et al. [Kim et al. 2012a] introduce the notion of fuzzy or probabilistic correspondence to facilitate a faceted shape retrieval system.

In contrast, Xu et al. [Xu et al. 2012a] introduce a *set evolution* technique, called fit-and-diverse, to create non-trivial 3D model

variations, starting from a set of structured models belonging to the same object class (see Figure 22). Instead of synthesizing new models one at a time, as in the classical modeling by example paradigm [Funkhouser et al. 2004; Kalogerakis et al. 2012] (see Figures 3 and 23), an entire set of models are evolved to produce novel shapes. Their method utilizes concepts from evolution theory to spawn new models while adapting to user intent in an interactive modeling setting. Specifically, in each evolution step, the user selects a subset of shapes considered to be fit which produce offsprings for the next generation. Thus, user preferences define the fitness for the evolution so that over time, the shape population is steered to keep good fitness.

To avoid overfitting which would result in an elite population with low diversity, two mechanisms are used to increase the diversity. First, a novel crossover operator is introduced that shuffles parts from different shapes in the collection to create new mutations. Second, a small percentage of individuals with lower fitness score are added to the next generation. As a result, the evolution produces inspiring, and sometimes unexpected, shapes. Finally, the part crossover operator works at the fine level of shape parts and allows new topologies to be created while still preserving the coarse-level functional structure of the input shapes.

Jain et al. [Jain et al. 2012] present a system to interpolate and blend between shapes from a database. They treat shapes as a collection of parts, which are held together by symmetry and contact constraints. The system then synthesizes new models by combining parts with comparable symmetry and contact structures.

Zheng et al. [Zheng et al. 2013] hypothesize that certain symmetric functional arrangements (i.e., part arrangements among symmetrically related parts) help capture object functions. Hence, preserving such arrangements in course of model synthesis, even when combining models from geometrically and topologically different sources (e.g., models from different collections) leads to functionally plausible models. Hence, they propose a purely geometric approach based on such substructures to match, replace, and position triplets of parts to create non-trivial, yet functionally plausible, model variations without requiring access to extensive training data (see Figure 24).

5.3 Structure-aware Shape Reconstruction

The GlobFit framework [Li et al. 2011a] regularizes locally fitted primitives based on extracted global mutual relations. The main assumption is that man-made engineering objects primarily consist of basic primitives, possibly repeated and globally aligned under common relations. Hence, they introduce an algorithm to directly couple the local and global aspects of the problem. The fitted primitives become the parts, while the relations are learned by testing for pairwise primitive relations. A set of feasible relations are extracted among the candidate relations, and then aligned to, while




Figure 23: Recombination of existing shapes based on a probabilistic model (green=training data, blue=synthesized result)[Kalogerakis et al. 2012].




Figure 24: Parts from different models, even those from different model families, can still be combined based on the consistency of relations among the parts, i.e., matching subgraphs in relation structures. This allows cross pollination of parts even across different model collections enabling functionally plausible model synthesis [Zheng et al. 2013].




Figure 25: Point set structuring has been proposed for surface reconstruction [Lafarge and Alliez 2013] even from very noisy and sparse input pointsets by directly recovering local structure (c.f., [Li et al. 2011b]) before surface reconstruction.

best fitting to the input data. Earlier, Li et al. [Li et al. 2011c] extract structure (in the form of translational symmetry) from photographs to improve model reconstruction from sparse and noisy LiDAR scans.

More recently, Lafarge and Alliez [Lafarge and Alliez 2013] reconstruct surfaces from point sets via a structure-preserving approach. The key idea is to use local planar components to consolidate point sets before actual reconstruction. The final surface is then obtained via a novel graph-cut problem formulated on the 3D Delaunay triangulation of the structured point set. The approach gracefully combines detected primitives with free-form parts of the inferred shape to produce robust and high-quality final models (see Figure 25).

The notion of context as structure information among co-occurring models was introduced in recent efforts by Fisher et al. [Fisher et al. 2011; Fisher et al. 2012], see also Figure 26. Structure has also been used recently for semantic model reconstruction from sparse, noisy, and often incomplete depth scans (e.g., Kinect scans). For example, Shen et al. [Shen et al. 2012] convert low-quality scans to high-quality 3D models using semantic labeled datasets (see Figure 27).

The algorithm runs bottom up and creates new structures by assembling existing labeled parts. Essentially the segmented input models become the parts, while the database suggests how they can be structurally linked. Alignment is performed to align the potential part arrangements to best explain the acquired data. In a parallel effort, Kim et al. [Kim et al. 2012b] propose a two phase approach: (i) a learning phase wherein they acquire 3D models of frequently occurring objects and capture their variability modes from multiple scans, and (ii) a recognition phase wherein from a single scan of a new environment, they identify previously seen objects but in novel poses and locations to support recognition at interactive rates.

5.4 Structure-aware Design Exploration

In the context of design exploration, a direct optimization approach (e.g., gradient descent) is rarely desirable. For example, designers are more interested in exploring the space of acceptable solutions, rather than being presented with a single (best) solution. Thus, the goal is to expose the space of acceptable solutions as characterized by the solutions that satisfy the structure relations (or constraints) within an allowed threshold margin.

Yang et al. [Yang et al. 2011] introduce the notion of *constrained meshes*, where the input is a single input mesh along with a set of non-linear constraints (in terms of the mesh vertices). The goal is to explore other meshes with the same connectivity, while respecting the prescribed constraints. They model the problem by mapping the meshes to points $\mathbf{x} \in \mathbb{R}^D$, where D is 3 times the number of deformable vertices. Each constraint defines a hypersurface in \mathbb{R}^D , and the intersection of all these hypersurfaces forms the corresponding shape space, or mesh manifold, \mathcal{M} . Thus, each energy E_i manifests as such a hypersurface. For example, face planarity leads to the planar quad (PQ) mesh manifold. They locally navigate in \mathcal{M} (actually within a tolerance band around \mathcal{M}) with the help of local approximations of \mathcal{M} . These are tangent spaces and quadratically parameterized surfaces having second order contact with \mathcal{M} . They observe that in many settings additional quality measures (e.g., aesthetics) are available to discriminate among the valid solutions. Thus, while any point of the shape space \mathcal{M} represents a valid constrained mesh, only certain parts of \mathcal{M} are desirable according to application specific quality measures, e.g., fairness of selected mesh polygons. They enable such an exploration with the help of appropriate energy functions and their second order approximations that are intrinsic to \mathcal{M} . Eigen-analysis of the intrinsic Hessians of energy functions turns out to be a highly effective tool for the identification of the good parts of the shape space for subsequent exploration with applications in design exploration, optimization, and handle driven deformation with boundary conditions,



Figure 26: Context-based model search result using graph kernels [Fisher et al. 2012].




Figure 27: Structure-aware 3D model reconstruction from Kinect data. The reconstruction is based on retrieving and then assembling parts taken from multiple example shapes in a small shape repository, while preserving the built-in structures of the examples [Shen et al. 2012].



Figure 28: Constrained meshes, i.e., the shape space of meshes with a set of prescribed constraints involving the free vertices, preserve non-local part relations while allowing the users to directly navigate in the implicitly defined constrained shape space [Yang et al. 2011].

while conforming to prescribed constraints (see Figure 28). The setup has subsequently been extended to handle curve-based model deformations [Zhao et al. 2012], and also deformations restricted to local modifications based on sparsity considerations [Deng et al. 2013].

Umetani et al. [Umetani et al. 2012] propose a guided modeling interface for exploration of valid shapes in the context of wood plank-based furniture. They propose an interactive design framework for efficient and intuitive exploration of geometrically and physically valid shapes. During geometric editing operation, the proposed system continuously visualizes the valid range of the parameter being edited. The planks (i.e., cuboids) become the parts, while the structures arise from the user-specified part layout. During manipulation, as one or more constraints (e.g., stability or toppling) are violated, the system generates multiple suggestions involving both discrete and continuous changes to restore validity. Thus, while the user focuses on the aesthetic aspects of the design, their computational framework helps to achieve physical realizability by providing active guidance to the user leading to fabrication- and physically-based form-finding possibility for conceptual design.

Whiting et al. [Whiting et al. 2012] also use structural constraints to guide form-finding for shapes, especially to masonry structures. They explore how geometric changes can lead to improved structural stability. The approach is based on a novel measure of structural soundness, which is then optimized by taking derivatives with respect to changing geometry. Specifically, geometric form is refined via a gradient descent while introducing other constraints like




Figure 29: Improving the stability of a building using structural optimization, thus enabling a novel form-finding possibility [Whiting et al. 2012].

volume preservation and uniform thickness (see Figure 29).

6 Conclusion

Inter-part relations within single shapes and across multiple shapes constitute structures of objects. In the context of man-made objects such structures capture characteristic non-local part relations and provide high-level information. Such part configurations along with their allowable variations arise out of various considerations: semantic factors, physical realizability, or fabrication convenience. Further, preserving such structures in course of any geometric manipulation help maintain the plausibility of the final geometric form. In recent years, there has been a flurry of research efforts to discover such structures from models, and subsequently preserve them in course of subsequent model manipulation. In this survey, we organize the various efforts in this emerging topic of structure-aware shape processing particularly focusing on extraction, representation, storage, and processing of structures in man-made shapes, with applications to various tasks including semantic reconstruction, scene understanding, intuitive deformation, plausible synthesis, and generally to design space exploration.

Future Challenges

We still consider structure-aware shape processing as a relatively young sub-field of computer graphics with many open issues and problems to address. We present a list of important open problem (in our opinion) to conclude this survey.

Hierarchical models. Most existing structure models represent flat arrangements of shape parts or features without a hierarchical organization, except for symmetry hierarchies [Martinet 2007; Wang et al. 2011]. However, symmetry hierarchies have so far been constructed by analyzing input shapes in isolation and the results do not ensure any consistency over a set of related shapes, limiting their utility. Since the way we perceive complex structures is generally thought of as being hierarchical, studies of hierarchical structure models for shape processing deserve more effort.

Functional shape analysis. The shift from low-level geometry processing to high-level structure-aware shape processing is posing new questions. Instead of focusing on measures that tell how similar two surface regions are geometrically, we are now interested in what makes two substructures functionally similar. Instead of producing shape variations which preserve local geometry details, we are now interested in how to maintain the same functionality during shape editing or creation. We expect functional shape analysis, i.e., how to properly map geometry to functionality, to draw more attention in research on geometry processing.

Scalability. To date, most works on structure-aware shape processing have only been demonstrated and evaluated on small to

medium-sized datasets over a handful of object categories. Popular object classes include those of chairs, airplanes, lamps, and other household items. 3D scene models typically contain up to twenty items. Hence, both in terms of data size and structural complexity, the scalability of the proposed analysis and modeling techniques has not been sufficiently investigated. In addition to scalability analysis, it is also clearly beneficial to develop large-scale benchmarks for various structure-aware processing tasks.

The data challenge. Much work covered in this survey is data-driven, where existing 3D shapes, possibly pre-analyzed or manually labeled, serve as either training data or examples for analysis or modeling. However, such approaches are always limited by the scant availability of quality 3D shapes and quality 3D shape collections. While the issue may resolve by itself as the number of digital 3D models continue to grow, one may look for alternative rich data sources. One possibility is to tap into the massive amount of image data, however, how image data can be utilized for structure-aware processing of 3D shapes is the challenge.

Unified structural representation. We have discussed a number of structure models in use so far and most of them were developed for specific applications. For example, in iWires [Gal et al. 2009], the model was built on feature curves and their relations and this was later extended for part-based analysis and manipulation via controllers [Zheng et al. 2011]. Data-driven modeling through reuse of existing shapes and their structures will be greatly facilitated if there were a unified structural representation. Such a representation requires the necessary formalism, compactness, and generality.

Conflicting constraints. As economical and accessible 3D data acquisition technologies, such as Kinect@home and LiDAR scans, propel the growth of 3D model collections, the quality of input data may become less unreliable and the structures extracted from such data may contain an increasing level of noise in the form of outliers or conflicting relations. Extracting a consistent set of relations from such inconsistent and conflicting relation sets will be another interesting and challenging research direction to pursue.

Author Background

Niloy J. Mitra is an associate professor (Reader) in the Department of Computer Science at University College London (UCL). Earlier, he co-founded the Geometric Modeling and Visualization Center (GMSV) center at KAUST. He received his Ph.D. degree and Masters in Electrical Engineering from Stanford University. His research interests include shape understanding, fabrication-aware design, geometric modeling, recreational art, and computer graphics in general. His work spans the area of symmetry detection, shape analysis, and their applications to semantic object manipulations. He received the ACM Siggraph Significant New Researcher Award in 2013 for his work on integrating form and function in 3D geometry.

Michael Wand is a senior scientist and junior research group leader at the Max-Planck Institut Informatik and Saarland University. He received his Ph.D. degree from Tübingen University in 2004 and his Diploma in Computer Science from Paderborn University in 2000. From 2005 to 2007, he was a postdoc at Stanford University. His research interests include statistical techniques for geometry processing and geometric correspondence problems.

Hao (Richard) Zhang directs the Graphics, Usability, and Visualization Lab at Simon Fraser University, Canada, where he is an associate professor in the School of Computing Science. He received his Ph.D. from the Dynamic Graphics Project (DGP), Department of Computer Science, University of Toronto in 2003 and M.Math. and B.Math. degrees from the University of Waterloo.

His research interests include geometry processing, shape analysis, and computer graphics. Currently, he serves on the editorial boards of Computer Graphics Forum and Graphical Models. He is a co-chair for SGP 2013. He was a winner of the Best Paper Award from SGP 2008 and the Most Cited Paper Award for the journal Computer-Aided Design in 2010.

Daniel Cohen-Or is a professor at the School of Computer Science. He received his B.Sc. cum laude in both mathematics and computer science (1985), an M.Sc. cum laude in computer science (1986) from Ben-Gurion University, and his Ph.D. degree from the Department of Computer Science (1991) at State University of New York at Stony Brook. He received the 2005 Eurographics Outstanding Technical Contributions Award. His research interests are in computer graphics, in particular, rendering and modeling techniques. His main interest right now is in few areas: image synthesis, motion and transformations, shapes and surfaces, and surface reconstruction.

Vladimir G. Kim received his PhD in the Computer Science Department at Princeton University and is current a postdoctoral scholar at Stanford University. His research interests include geometry processing and analysis of shapes and collections of 3D models. He received his B.A. degree in Mathematics and Computer Science from Simon Fraser University in 2008. Vladimir is a recipient of the Siebel Scholarship and the NSERC Postgraduate Scholarship. He is also on the International Program Committee for SGP 2013.

Qi-xing Huang is a postdoc researcher at Stanford University. He earned his Ph.D. from the Department of Computer Science at Stanford University in 2011. He obtained both MS and BS degrees in Computer Science from Tsinghua University in 2005 and 2002, respectively. His research interests include data-driven geometry processing and co-analysis of shapes and collections of 3D models using convex optimization techniques. He was a winner of the Best Paper Award from SGP 2013 and the Most Cited Paper Award for the journal Computer-Aided Geometric Design in 2011 and 2012.

References

- ALHASHIM, I., ZHANG, H., AND LIU, L. 2012. Detail-replicating shape stretching. *The Visual Computer* 28, 12, 1153–1166.
- ALIAGA, D., ROSEN, P., AND BEKINS, D. 2007. Style grammars for interactive visualization of architecture. *IEEE Trans. Vis. Comp. Graph.* 13, 4, 786–797.
- ALLEN, B., CURELLS, B., AND POPOVIĆ, Z. 2003. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 3, 587–594.
- ALT, H., MEHLHORN, K., WAGENER, H., AND WELZL, E. 1988. Congruence, similarity and symmetries of geometric objects. *Discrete Comput. Geom.* 3, 3, 237–256.
- ANGUELOV, D., TASKAR, B., CHATALBASHEV, V., KOLLER, D., GUPTA, D., HEITZ, G., AND NG, A. 2005. Discriminative learning of markov random fields for segmentation of 3d scan data. In *CVPR*.
- ATTENE, M., ROBBIANO, F., SPAGNUOLO, M., AND FALCIDIENO, B. 2009. Characterization of 3d shape parts for semantic annotation. *Computer Aided Design* 41, 10, 756–763.
- BAO, F., SCHWARZ, M., AND WONKA, P. 2012. Procedural facade variations from a single layout. *ACM Trans. Graph.*
- BHAT, P., INGRAM, S., AND TURK, G. 2004. Geometric texture synthesis by example. In *Symp. Geometry Processing*, 41–44.

- BLANZ, V., AND VETTER, T. 1999. A morphable model for the synthesis of 3d faces. In *Computer Graphics (Proc. SIGGRAPH)*, 187–194.
- BOKELOH, M., BERNER, A., WAND, M., SEIDEL, H.-P., AND SCHILLING, A. 2009. Symmetry detection using line features. *Computer Graphics Forum* 28, 2, 697–706.
- BOKELOH, M., WAND, M., AND SEIDEL, H.-P. 2010. A connection between partial symmetry and inverse procedural modeling. *ACM Trans. Graph.* 29 (July), 104:1–104:10.
- BOKELOH, M., WAND, M., KOLTUN, V., AND SEIDEL, H.-P. 2011. Pattern-aware shape deformation using sliding dockers. *ACM Trans. Graph.* 30 (Dec.), 123:1–123:10.
- BOKELOH, M., WAND, M., SEIDEL, H.-P., AND KOLTUN, V. 2012. An algebraic model for parameterized shape editing. *ACM Transactions on Graphics* 31, 4, 78:1–78:10.
- BOTSCH, M., AND SORKINE, O. 2008. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics* 14, 1, 213–230.
- CHAUDHURI, S., KALOGERAKIS, E., GUIBAS, L., AND KOLTUN, V. 2011. Probabilistic reasoning for assembly-based 3D modeling. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4, 35:1–35:10.
- CHEN, L., AND MENG, X. 2009. Anisotropic resizing of model with geometric textures. In *Conf. on Geometric and Physical Modeling (SPM)*, ACM, New York, NY, USA, 289–294.
- COQUILLART, S. 1990. Extended free-form deformation: a sculpturing tool for 3d geometric modeling. In *Proc. Siggraph*.
- DENG, B., BOUAZIZ, S., DEUSS, M., ZHANG, J., SCHWARTZBURG, Y., AND PAULY, M. 2013. Exploring local modifications for constrained meshes. *Computer Graphics Forum (Proceedings of Eurographics 2013)* 32, 2, to appear.
- DUDA, R. O., HART, P. E., AND STORK, D. G. 2000. *Pattern Classification (2nd Edition)*. Wiley-Interscience.
- EFROS, A. A., AND LEUNG, T. K. 1999. Texture synthesis by non-parametric sampling. In *Proc. Int. Conf. Comp. Vision*.
- FISHER, M., SAVVA, M., AND HANRAHAN, P. 2011. Characterizing structural relationships in scenes using graph kernels. *ACM Trans. Graph.* 30, 4, 34:1–34:12.
- FISHER, M., RITCHIE, D., SAVVA, M., FUNKHOUSER, T., AND HANRAHAN, P. 2012. Example-based synthesis of 3d object arrangements. *ACM Transactions on Graphics (SIGGRAPH Asia)* 31, 6, 135:1–135:11.
- FU, H., COHEN-OR, D., DROR, G., AND SHEFFER, A. 2008. Upright orientation of man-made objects. *ACM Trans. Graph.* 27, 3 (Aug.), 42:1–42:7.
- FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. Graph.* 23, 3.
- GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph.* 25, 1, 130–150.
- GAL, R., SORKINE, O., MITRA, N., AND COHEN-OR, D. 2009. iWires: An Analyze-and-Edit Approach to Shape Manipulation. *ACM Trans. Graph.* 28, 3, 33:1–33:10.
- GELFAND, N., AND GUIBAS, L. 2004. Shape segmentation using local slippage analysis. In *Proc. Symp. Geometry Processing*, 214–223.
- GOLOVINSKIY, A., AND FUNKHOUSER, T. 2009. Consistent segmentation of 3D models. *Computers & Graphics (Proc. of SMI)* 33, 3, 262–269.
- HABBECKE, M., AND KOBELT, L. 2012. Linear analysis of non-linear constraints for interactive geometric modeling. *Computer Graphics Forum* 31, 2.
- HART, J., AND FLYNN, O. C. P. 1997. Similarity hashing: A computer vision solution to the inverse problem of linear fractals. *Fractals* 5, 35–50.
- HERTZMANN, A., JACOBS, C. E., OLIVER, N., CURLESS, B., AND SALESIN, D. H. 2001. Image analogies. In *Proc. Siggraph 2001*, 327–340.
- HOFFMANN, C. M., AND JOAN-ARINYO, R. 2005. A brief on constraint solving. *Computer-Aided Design and Applications* 2, 5, 655–663.
- HU, R., FAN, L., AND LIU, L. 2012. Co-segmentation of 3d shapes via subspace clustering. *Computer Graphics Forum (Proceedings of SGP)* 31, 5, 1703–1713.
- HUANG, Q., MECH, R., AND CARR, N. 2009. Optimizing structure preserving embedded deformation for resizing images and vector art. vol. 28, 1887–1896.
- HUANG, Q., KOLTUN, V., AND GUIBAS, L. 2011. Joint shape segmentation with linear programming. *ACM Trans. Graph.* 30, 125:1–125:12.
- IGARASHI, T., MATSUOKA, S., AND TANAKA, H. 1999. Teddy: A sketching interface for 3d freeform design. In *Computer Graphics (Proc. ACM SIGGRAPH)*, 409–416.
- IJIRI, T., OWADA, S., AND IGARASHI, T. 2006. The sketch I-system: Global control of tree modeling using free-form strokes. In *Smart Graphics*, 138–146.
- IJIRI, T., MĚCH, R., IGARASHI, T., AND MILLER, G. 2008. An example-based procedural system for element arrangement. *Computer Graphics Forum* 27, 3, 429–436.
- JAIN, A., THORMÄHLEN, T., RITSCHER, T., AND SEIDEL, H.-P. 2012. Exploring shape variations by 3d-model decomposition and part-based recombination. *Comp. Graph. Forum (Proc. Eurographics 2012)* 31, 2, 631–640.
- KALOGERAKIS, E., HERTZMANN, A., AND SINGH, K. 2010. Learning 3d mesh segmentation and labeling. *ACM Trans. Graph.* 29, 3.
- KALOGERAKIS, E., CHAUDHURI, S., KOLLER, D., AND KOLTUN, V. 2012. A probabilistic model for component-based shape synthesis. *ACM Transactions on Graphics* 31, 4, 55:1–55:11.
- KALOJANOV, J., BOKELOH, M., WAND, M., GUIBAS, L., SEIDEL, H.-P., AND SLUSALLEK, P. 2012. Microtiles: Extracting building blocks from correspondences. John Wiley & Sons, Inc., New York, NY, USA, vol. 31, 1597–1606.
- KENT, J. T., MARDIA, K. V., WEST, J. M., AND JT, L. L. 1996. Ridge curves and shape analysis. In *In The British Machine Vision Conference*, 43–52.

- KILIAN, M., MITRA, N. J., AND POTTMANN, H. 2007. Geometric modeling in shape space. *ACM Transactions on Graphics (SIGGRAPH)* 26, 3, #64, 1–8.
- KIM, V. G., LI, W., MITRA, N., DIVERDI, S., AND FUNKHOUSER, T. 2012. Exploring collections of 3d models using fuzzy correspondences. *Trans. on Graphics (Proc. of SIGGRAPH 2012)* 31, 4, 54:1–54:11.
- KIM, Y. M., MITRA, N. J., YAN, D.-M., AND GUIBAS, L. 2012. Acquiring 3d indoor environments with variability and repetition. *ACM Transactions on Graphics (SIGGRAPH Asia)* 31, 6, 138:1–138:11.
- KOLLER, D., AND FRIEDMAN, N. 2009. *Probabilistic Graphical Models*. MIT Press.
- KRAEVOY, V., JULIUS, D., AND SHEFFER, A. 2007. Shuffler: Modeling with interchangeable parts. In *Pacific Graphics 2007*.
- KRAEVOY, V., SHEFFER, A., SHAMIR, A., AND COHEN-OR, D. 2008. Non-homogeneous resizing of complex models. *ACM Trans. Graph.* 27, 5, 1–9.
- KUMAR, S., AND HEBERT, M. 2003. Discriminative random fields: A discriminative framework for contextual interaction in classification. In *ICCV*.
- KWATRA, V., SCHÖDL, A., ESSA, I., TURK, G., AND BOBICK, A. 2003. Graphcut textures: image and video synthesis using graph cuts. *ACM Trans. Graph.* 22, 3, 277–286.
- KWATRA, V., ESSA, I., BOBICK, A., AND KWATRA, N. 2005. Texture optimization for example-based synthesis. *ACM Trans. Graph.* 24, 3, 795–802.
- LAFARGE, F., AND ALLIEZ, P. 2013. Surface reconstruction through point set structuring. In *Eurographics*. to appear.
- LAGAE, A., DUMONT, O., AND DUTRÉ, P. 2005. Geometry synthesis by example. In *Conf. Shape Modeling and Applications*.
- LAI, Y.-K., HU, S.-M., GU, D. X., AND MARTIN, R. R. 2005. Geometric texture synthesis and transfer via geometry images. In *Symp. Solid and Physical Modeling*, 15–26.
- LI, Y., WU, X., CHRYSANTHOU, Y., SHARF, A., COHEN-OR, D., AND MITRA, N. J. 2011. Globfit: Consistently fitting primitives by discovering global relations. *ACM Transactions on Graphics* 30, 4, 52:1–52:12.
- LI, Y., WU, X., CHRYSANTHOU, Y., SHARF, A., COHEN-OR, D., AND MITRA, N. J. 2011. Globfit: Consistently fitting primitives by discovering global relations. *ACM Transactions on Graphics* 30, 4, 52:1–52:12.
- LI, Y., ZHENG, Q., SHARF, A., COHEN-OR, D., CHEN, B., AND MITRA, N. J. 2011. 2d-3d fusion for layer decomposition of urban facades. In *ICCV*.
- LIN, J., COHEN-OR, D., ZHANG, H. R., LIANG, C., SHARF, A., DEUSSEN, O., AND CHEN, B. 2011. Structure-preserving retargeting of irregular 3d architecture. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2011)* 30, 6 (dec), 183:1–183:10.
- LIPP, M., WONKA, P., AND WIMMER, M. 2008. Interactive visual editing of grammars for procedural architecture. *ACM Trans. Graph.* 27, 3, 1–10.
- LIU, Y., COLLINS, R., AND TSIN, Y. 2004. A computational model for periodic pattern perception based on frieze and wall-paper groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 1 (March), 354 – 371.
- LIU, Y., HEL-OR, H., KAPLAN, C. S., AND GOOL, L. V. 2010. Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision* 5, 1–195.
- LOY, G., AND EKLUNDH, J. 2006. Detecting symmetry and symmetric constellations of features. In *Proc. Europ. Conf. Computer Vision*, 508–521.
- MARTINET, A., SOLER, C., HOLZSCHUCH, N., AND SILLION, F. 2006. Accurate detection of symmetries in 3d shapes. *ACM Trans. Graph.* 25, 2, 439 – 464.
- MARTINET, A. 2007. *Structuring 3D Geometry based on Symmetry and Instancing Information*. PhD thesis, INP Grenoble.
- MERRELL, P., AND MANOCHA, D. 2008. Continuous model synthesis. *ACM Trans. Graph.* 27, 5, 1–7.
- MERRELL, P., AND MANOCHA, D. 2009. Constraint-based model synthesis. In *SPM '09: 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling*, ACM, New York, NY, USA, 101–111.
- MERRELL, P. 2007. Example-based model synthesis. In *Symp. Interactive 3D Graphics and Games*, 105–112.
- MITRA, N. J., AND PAULY, M. 2008. Symmetry for architectural design. In *Advances in Architectural Geometry*, 13–16.
- MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2006. Partial and approximate symmetry detection for 3d geometry. *ACM Trans. Graph.* 25, 3, 560–568.
- MITRA, N. J., GUIBAS, L., AND PAULY, M. 2007. Symmetrization. vol. 26, Article 63.
- MITRA, N. J., YANG, Y.-L., YAN, D.-M., LI, W., AND AGRAWALA, M. 2010. Illustrating how mechanical assemblies work. *ACM Transactions on Graphics* 29, 3.
- MITRA, N. J., PAULY, M., WAND, M., AND CEYLAN, D., 2012. Symmetry in 3d geometry: Extraction and applications. Eurographics State-of-the-art Report (STAR).
- MITRA, N. J., YANG, Y.-L., YAN, D.-M., LI, W., AND AGRAWALA, M. 2013. Illustrating how mechanical assemblies work. *ACM CACM Research Highlight*.
- MÜLLER, P., WONKA, P., HAEGLER, S., ULMER, A., AND GOOL, L. V. 2006. Procedural modeling of buildings. *ACM Trans. Graph.* 25, 3, 614–623.
- NGUYEN, M. X., YUAN, X., AND CHEN, B. 2005. Geometry completion and detail generation by texture synthesis. *The Visual Computer* 21, 9–10, 669–678.
- OHTAKE, Y., BELYAEV, A., AND SEIDEL, H.-P. 2004. Ridge-valley lines on meshes via implicit surface fitting. In *Proc. Siggraph*, 609–612.
- OVSJANIKOV, M., SUN, J., AND GUIBAS, L. 2008. Global intrinsic symmetries of shapes. In *Eurographics Symposium on Geometry Processing (SGP)*.
- OVSJANIKOV, M., LI, W., GUIBAS, L., AND MITRA, N. J. 2011. Exploration of continuous variability in collections of 3d shapes. *ACM Transactions on Graphics* 30, 4, 33:1–33:10.
- PARISH, Y. I. H., AND MÜLLER, P. 2001. Procedural modeling of cities. In *Proc. Siggraph 2001*, 301–308.

- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. 2008. Discovering structural regularity in 3D geometry. *ACM Trans. Graph.* 27, 3.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Trans. Graph.* 25, 3.
- RAVIV, D., BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2010. Full and partial symmetries of non-rigid shapes. *Intl. Journal of Computer Vision (IJCV)* 89, 1 (August), 18–39.
- SEDERBERG, T. W., AND PARRY, S. R. 1986. Free-form deformation of solid geometric models. In *Proc. Siggraph*, 151–160.
- SHAMIR, A. 2008. A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6, 1539–1556.
- SHARF, A., ALEXA, M., AND COHEN-OR, D. 2004. Context-based surface completion. *ACM Trans. Graph.* 23, 3, 878–887.
- SHAWE-TAYLOR, J., AND CRISTIANINI, N. 2004. *Kernel Methods for Pattern Analysis*. Cambridge University Press.
- SHEN, C.-H., FU, H., CHEN, K., AND HU, S.-M. 2012. Structure recovery by part assembly. *ACM Transactions on Graphics (SIGGRAPH Asia)* 31, 6, 180:1–180:11.
- SIDI, O., VAN KAICK, O., KLEIMAN, Y., ZHANG, H., AND COHEN-OR, D. 2011. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 30, 6, 126:1–126:10.
- SIMAKOV, D., CASPI, Y., SHECHTMAN, E., AND IRANI, M. 2008. Summarizing visual data using bidirectional similarity. In *IEEE Conf. Comp. Vision and Pattern Recognition (CVPR)*, 1–8.
- SIMARI, P., KALOGERAKIS, E., AND SINGH, K. 2006. Folding meshes: hierarchical mesh segmentation based on planar symmetry. In *Proc. Symp. Geometry Processing*, 111–119.
- SINGH, K., AND EUGENE, F. 1998. Wires: a geometric deformation technique. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '98, 405–414.
- SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *Proceedings of Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, 109–116.
- STÁVA, O., BENES, B., MĚCH, R., ALIAGA, D., AND KRIŠTOF, P. 2010. Inverse procedural modeling by automatic generation of l-systems. *Computer Graphics Forum* 29, 2, 665–674.
- STINY, G., AND GIPS, J. 1971. Shape grammars and the generative specification of painting and sculpture. In *IFIP Congress 71*.
- SUNKEL, M., JANSEN, S., WAND, M., AND SEIDEL, H.-P. 2013. A correlated parts model for object detection in large 3d scans. *Computer Graphics Forum (Special Issue of Eurographics)*. to appear.
- TALTON, J. O., LOU, Y., LESSER, S., DUKE, J., MĚCH, R., AND KOLTUN, V. 2011. Metropolis procedural modeling. *ACM Trans. Graph.* 30 (April), 11:1–11:14.
- TALTON, J., YANG, L., KUMAR, R., LIM, M., GOODMAN, N. D., AND MĚCH, R. 2012. Learning design patterns with bayesian grammar induction. In *Proceedings of the 25th ACM Symposium on User Interface Software and Technology*, ACM, New York, NY, USA, UIST '12, 63–74.
- TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *Proc. SIGGRAPH '87*, ACM, New York, NY, USA, 205–214.
- THOMPSON, D. W. 1992. *On Growth and Form*. Dover reprint of 1942 2nd ed. (1st ed., 1917).
- THRUN, S., AND WEGBREIT, B. 2005. Shape from symmetry. In *Proc. Int. Conf. Computer Vision*.
- UMETANI, N., IGARASHI, T., AND MITRA, N. J. 2012. Guided exploration of physically valid shapes for furniture design. *ACM Transactions on Graphics* 31, 4, 86:1–86:11.
- VAN KAICK, O., TAGLIASACCHI, A., SIDI, O., ZHANG, H., COHEN-OR, D., WOLF, L., AND HAMARNEH, G. 2011. Prior knowledge for shape correspondence. *Computer Graphics Forum (Special Issue of Eurographics)* 30, 2, 553–562.
- VAN KAICK, O., ZHANG, H., HAMARNEH, G., AND COHEN-OR, D. 2011. A survey on shape correspondence. *Computer Graphics Forum* 30, 6, 1681–1707.
- VANEGAS, C., DORADO, I. G., ALIAGA, D., BENES, B., AND WADDELL, P. 2012. Inverse design of urban procedural models. *ACM Transactions on Graphics* 31, 6, 168:1–168:11.
- WANG, Y., XU, K., LI, J., ZHANG, H., SHAMIR, A., LIU, L., CHENG, Z., AND XIONG, Y. 2011. Symmetry hierarchy of man-made objects. vol. 30, 287–296.
- WANG, Y., ASAFI, S., VAN KAICK, O., ZHANG, H., COHEN-OR, D., AND CHEN, B. 2012. Active co-analysis of a set of shapes. *ACM Transactions on Graphics* 31, 6.
- WEI, L., AND LEVOY, M. 2000. Fast texture synthesis using tree-structured vector quantization. In *Proc. Siggraph 2000*, 479–488.
- WEI, L., HAN, J., ZHOU, K., BAO, H., GUO, B., AND SHUM, H. 2008. Inverse texture synthesis. *ACM Trans. Graph.* 27, 3, 52:1–52:9.
- WERTHEIMER, M. 1923. Untersuchungen zur lehre von der gestalt. ii. *Psychologische Forschung* 4, 301–350.
- WHITING, E., SHIN, H., WANG, R., OCHSENDORF, J., AND DURAND, F. 2012. Structural optimization of 3d masonry buildings. *ACM Transactions on Graphics* 31, 6.
- WONKA, P., WIMMER, M., SILLION, F., AND RIBARSKY, W. 2003. Instant architecture. *ACM Trans. Graph.* 22, 3, 669–677.
- XU, K., ZHANG, H., TAGLIASACCHI, A., LIU, L., LI, G., MENG, M., AND XIONG, Y. 2009. Partial intrinsic reflectional symmetry of 3d shapes. *ACM Transactions on Graphics (Special Issue of SIGGRAPH Asia)* 28, 5, 138:1–138:10.
- 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. Graph.* 28, 3, 1–9.
- XU, K., LI, H., ZHANG, H., COHEN-OR, D., XIONG, Y., AND CHENG, Z. 2010. Style-content separation by anisotropic part scales. *ACM Transactions on Graphics, (Proc. of SIGGRAPH Asia 2010)* 29, 5.
- XU, K., ZHENG, H., ZHANG, H., COHEN-OR, D., LIU, L., AND XIONG, Y. 2011. Photo-inspired model-driven 3d object modeling. *ACM Transactions on Graphics* 30, 4, 80:1–10.
- XU, K., ZHANG, H., COHEN-OR, D., AND CHEN, B. 2012. Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM*

- Transactions on Graphics (Special Issue of SIGGRAPH)* 31, 4, 57:1–57:10.
- XU, K., ZHANG, H., JIANG, W., DYER, R., CHENG, Z., LIU, L., AND CHEN, B. 2012. Multi-scale partial intrinsic symmetry detection. *ACM Transactions on Graphics (Special Issue of SIGGRAPH Asia)* 31, 6, 181:1–181:10.
- YANG, Y.-L., YANG, Y.-J., POTTMANN, H., AND MITRA, N. J. 2011. Shape space exploration of constrained meshes. *ACM Trans. Graph.* 30, 6 (Dec.), 124:1–124:12.
- YEH, Y.-T., AND MĚCH, R. 2009. Detecting symmetries and curvilinear arrangements in vector art. *Computer Graphics Forum* 28, 2, 707–716.
- YEH, Y.-T., BREEDEN, K., YANG, L., FISHER, M., , AND HANRAHAN, P. 2013. Synthesis of tiled patterns using factor graphs. *ACM Transactions on Graphics*, to appear.
- YUMER, M., AND KARA, L. 2012. Co-abstraction of shape collections. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2012)* 31, 158:1–158:11.
- ZELEZNIK, R. C., HERNDON, K. P., AND HUGHES, J. F. 2007. Sketch: an interface for sketching 3d scenes. In *ACM SIGGRAPH 2007 courses*, ACM, New York, NY, USA, SIGGRAPH '07.
- ZELINKA, S., AND GARLAND, M. 2004. Similarity-based surface modelling using geodesic fans. In *Proc. Symp. Geometry processing*, 204–213.
- ZELINKA, S., AND GARLAND, M. 2006. Surfacing by numbers. In *Graphics Interface*, 107–113.
- ZHAO, X., TANG, C.-C., YANG, Y.-L., POTTMANN, H., AND MITRA, N. J. 2012. Intuitive design exploration of constrained meshes. In *Advances in Architectural Geometry*.
- ZHENG, Y., FU, H., COHEN-OR, D., AU, O. K.-C., AND TAI, C.-L. 2011. Component-wise controllers for structure-preserving shape manipulation. In *Computer Graphics Forum (Proc. Eurographics)*.
- ZHENG, Y., CHEN, X., CHENG, M.-M., ZHOU, K., HU, S.-M., AND MITRA, N. J. 2012. Interactive images: Cuboid proxies for smart image manipulation. *ACM Transactions on Graphics* 31, 4, 99:1–99:11.
- ZHENG, Y., COHENOR, D., AND MITRA, N. J. 2013. Functional substructures for part compatibility. *Computer Graphics Forum*, to appear.
- ZHOU, K., HUANG, X., WANG, X., TONG, Y., DESBRUN, M., AND BAINING GUO, H.-Y. S. 2006. Mesh quilting for geometric texture synthesis. *ACM Trans. Graph.* 25, 3, 690–697.
- ZHU, L., XU, W., SNYDER, J., LIU, Y., WANG, G., AND GUO, B. 2012. Motion-guided mechanical toy modeling. *ACM Trans. Graph.* 31, 6.