



# Data-Driven Shape Analysis and Processing

Kai Xu<sup>1,2</sup>, Vladimir G. Kim<sup>3,4</sup>, Qixing Huang<sup>5</sup> and Evangelos Kalogerakis<sup>6</sup>

<sup>1</sup>National University of Defense Technology, Changsha, Hunan, China  
kaixu@nudt.edu.cn

<sup>2</sup>Shenzhen VisuCA Key Lab/SIAT, Shenzhen, China

<sup>3</sup>Department of Computer Science, Stanford University, Stanford, CA, USA

<sup>4</sup>Adobe Research, Seattle, WA, USA  
vokim@adobe.com

<sup>5</sup>Toyota Technological Institute at Chicago, Chicago, IL, USA  
huangqx@ttic.edu

<sup>6</sup>College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA  
kalo@cs.umass.edu

---

## Abstract

*Data-driven methods serve an increasingly important role in discovering geometric, structural and semantic relationships between shapes. In contrast to traditional approaches that process shapes in isolation of each other, data-driven methods aggregate information from 3D model collections to improve the analysis, modelling and editing of shapes. Data-driven methods are also able to learn computational models that reason about properties and relationships of shapes without relying on hard-coded rules or explicitly programmed instructions. Through reviewing the literature, we provide an overview of the main concepts and components of these methods, as well as discuss their application to classification, segmentation, matching, reconstruction, modelling and exploration, as well as scene analysis and synthesis. We conclude our report with ideas that can inspire future research in data-driven shape analysis and processing.*

**Keywords:** Shape analysis, shape processing, shape modeling, data-driven approach, machine learning

**ACM CCS:** I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling

---

## 1. Introduction

As the availability of 3D data increases, due to the developments in both 3D sensing technology as well as 3D modelling software, data-driven approaches become increasingly applicable and useful to 3D shape processing. In contrast to traditional approaches [LZ11], data-driven methods look beyond single objects, instead analysing sets of shapes jointly to extract meaningful mappings and correlations between them. In addition, these methods are able to learn from data computational models that effectively reason about properties and relationships of shapes without relying on hard-coded rules or explicitly programmed instructions. Leveraging shared information across multiple objects, data-driven methods are able to facilitate high-level shape understanding through discovering geometric and structural patterns among collections of shapes, patterns which serve as strong priors in various geometry processing applications.

The idea of utilizing data to support geometry processing has been exploited and practiced for many years. However, most existing works based on this idea are confined to the example-based paradigm, mostly leveraging only one core concept of data-driven techniques—*information transfer*. Typically, the input to these problems includes one or multiple exemplar shapes with some prescribed or pre-computed information of interest, and a target shape that needs to be analysed or processed. These techniques usually establish a *correlation* between the source and the target shapes and transfer the interesting information from the source to the target. The applications of such approaches include a variety of methods in shape analysis (e.g. [SY07]) and shape synthesis (e.g. [Mer07, MHS\*14]).

As the availability of 3D data increases, several new concepts in data-driven methods are emerging, opening space for new developments in shape analysis and content creation. First, the rich

variability of 3D content in existing shape repositories makes it possible to directly reuse the shapes or parts for constructing new 3D models [FKS\*04]. *Content reuse* for 3D modelling is perhaps the most straightforward application of big 3D geometric data, providing a promising approach to address the challenging 3D content creation problem. Second, high-level shape understanding can benefit from co-analysing collections of shapes. Several analysis tools demonstrate that shape analysis is more reliable if it is supported by observing certain attributes across a set of semantically related shapes instead of just focusing on a single object. *Co-analysis* requires a critical step of finding correlations between multiple shapes in the input set, which is substantially different from building pairwise correlation. A key concept to co-analysis is the *consistency* of the correlations across the entire set, which has both semantic [KHS10, SvKK\*11, WAvK\*12] and mathematical [HG13] justifications. Third, aside from analysing patterns from a set of shapes, it is also possible to endorse a subset of the shapes with some semantic information (e.g. part labelling), which can be propagated to the other shapes through learned mappings. This information propagation evolves the concept of *knowledge transfer* between shapes.

**Relation to knowledge-driven shape processing.** Prior to the emergence of data-driven techniques, high-level shape understanding and modelling was usually achieved with knowledge-driven methods. In the knowledge-driven paradigm, geometric and structural patterns are extracted and interpreted with the help of explicit rules or hand-crafted parameters. Such examples include heuristics-based shape segmentation [Sha08] and procedural shape modelling [MWH\*06]. Although these approaches have certain empirical success, they exhibit several inherent limitations. First, it is extremely difficult to hard-code explicit rules and heuristics that can handle the enormous geometric and structural variability of 3D shapes and scenes in general. As a result, knowledge-driven approaches are often hard to generalize well to large and diverse shape collections. Another issue is that non-experts find it difficult to interact with knowledge-driven techniques that require as input ‘low-level’ geometric parameters or instructions.

In contrast to knowledge-driven methods, data-driven techniques learn representations and parameters from data. They usually do not depend on hard-coded prior knowledge, and consequently do not rely on hand-crafted parameters, making these techniques more data-adaptive and thus lead to significantly improved performance in many practical settings. The success of data-driven approaches, backed by machine learning techniques, heavily relies on the accessibility of large data collections. We have witnessed the successful performance improvement of machine learning algorithms by increasing the training set size [BB01]. In light of this, the recent developments in 3D modelling tools and acquisition techniques for 3D geometry, as well as availability of large repositories of 3D shapes (e.g. Trimble 3D Warehouse, Yobi3D, etc.), offer great opportunities for developing data-driven approaches for 3D shape analysis and processing.

**Relation to structure-aware shape processing.** This report is closely related to the recent survey on ‘structure-aware shape processing’ by Mitra and co-workers [MWZ\*14], which concentrates on techniques for structural analysis of 3D shapes, as well as high-level shape processing guided by structure-preservation. In that survey, shape structure is defined as the arrangement and relations

between shape parts, which is analysed through identifying shape parts, part parameters and part relations. Each of the three can be determined through manual assignment, pre-defined model fitting and data-driven learning.

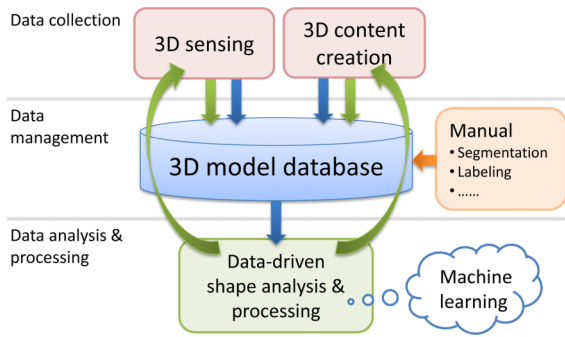
In contrast, our report takes a very different perspective—we focus on how the increasing availability of geometric data has changed the field of shape analysis and processing. In particular, we want to highlight several key distinctions: First, data-driven shape processing goes beyond structure analysis. For example, leveraging large shape collections may benefit a wider variety of problems in shape understanding and processing, such as parametric modelling of shape space [ACP03], hypothesis generation for object and scene understanding [ZSSS13, SLH12] and information transfer between multi-modal data [WGW\*13, SHM\*14]. Data-driven shape processing may also exploit the data-centred techniques in machine learning such as sparse representation [RR13] and feature learning [HOT06, Ben09, YN10, KSH12], which are not pre-conditioned on any domain-specific or structural prior beyond raw data. Second, even within the realm of structure-aware shape processing, data-driven approaches are arguably becoming dominant due to their theoretical and practical advantages, as well as the availability of large shape repositories and recent developments in machine learning.

**Vision and motivation.** With the emergence of ‘big data’, many scientific disciplines have shifted their focus to data-driven techniques. Although 3D geometry data are still far from being as ubiquitous as some other data formats (e.g. photographs), the rapidly growing number of 3D models, the recent developments in fusing 2D and 3D data, and the invention of commodity depth sensors, have made the era of ‘big 3D data’ more promising than ever. At the same time, we expect data-driven approaches to take one of the leading roles in the reconstruction and understanding of acquired 3D data, as well as the synthesis of new shapes. Data-driven geometry processing will close the loop starting from acquisition, analysis and processing all the way to the generation of 3D shapes (see Figure 1), and will be a key tool for manipulating big visual data.

Recent years have witnessed a rapid development of data-driven geometry processing algorithms, both in the computer graphics and computer vision communities. Given the research efforts and wide interests in the subject, we believe many researchers would benefit from a comprehensive and systematic survey. We also hope such a survey can stimulate new theories, problems and applications.

**Organization.** This survey is organized as follows. Section 2 gives a high-level overview of data-driven approaches and classifies data-driven methods with respect to their application domains. This section also provides two representative examples for the readers to understand the general workflow of data-driven geometry processing. The sections following survey the various data-driven shape processing problems in detail, and try to correlate the different methods through comparisons in various aspects. Finally, we conclude our survey by discussing a few key problems involved in designing a data-driven method for shape processing, listing a set of open challenges in this direction, as well as providing a vision on future research.

**Accompanying online resources.** In order to assist the reader in learning and leveraging the basic algorithms, we provide an online wikipedia [XKHK14], which collects tools and source code,



**Figure 1:** Data-driven shape processing and modelling provides a promising solution to the development of ‘big 3D data’. The two major ways of 3D data generation, 3D sensing and 3D content creation, populate 3D databases with fast growing amount of 3D models. The database models are sparsely augmented with manual segmentation and labelling, as well as reasonably organized, to support data-driven shape analysis and processing, based on, e.g. machine learning techniques. The learned knowledge can in turn support efficient 3D reconstruction and 3D content creation, during which the knowledge can be transferred to the newly generated data. Such 3D data with semantic information can be included into the database to enrich it and facilitate further data-driven applications.

together with benchmark data for typical problems and applications of data-driven shape processing. This page will also maintain links and data mining tools for obtaining large data collections of shapes and scenes. This website could serve as a starting point for those who are conducting research in this direction. We also expect it to benefit a wide spectrum of researchers from related fields.

## 2. Overview

In this section, we provide a high-level overview of the main components and steps of data-driven approaches for processing 3D shapes and scenes. Although the pipeline of these methods vary significantly depending on their particular applications and goals, a number of components tend to be common: the input data collection and processing, data representations and feature extraction, as well as learning and inference. *Representation, learning and inference* are critical components of machine learning approaches in general [KF09]. In the case of shape and scene processing, each of these components poses several interesting and unique problems when dealing with 3D geometric data. These problems have greatly motivated the research on data-driven geometry processing, and in turn have brought new challenges to the computer vision and machine learning communities, as reflected by the increased interest in 3D visual data from these fields. Below, we discuss particular characteristics and challenges of data-driven 3D shape and scene processing algorithms. Figure 2 provides a schematic overview of the most common components of these algorithms.

### 2.1. 3D data collection

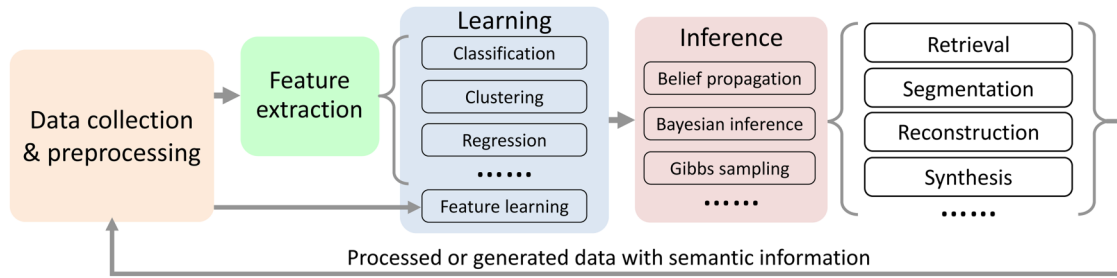
**Shape representation.** A main component of data-driven approaches for shape and scene processing is data collection, where

the goal is acquire a number of 3D shapes and scenes depending on the application. When shapes and scenes are captured with scanners or depth sensors, their initial representation is in the form of *range data* or *unorganized point clouds*. Several data-driven methods for reconstruction, segmentation and recognition directly work on these representations and do not require any further processing. On the other hand, online repositories, such as the Trimble 3D Warehouse, contain millions of shapes and scenes that are represented as *polygon meshes*. A large number of data-driven techniques are designed to handle complete shapes in the form of polygon meshes created by 3D modelling tools or reconstructed from point clouds. Choosing which representation to use depends on the application. For example, data-driven reconstruction techniques aim for generating complete shapes and scenes from noisy point clouds with missing data. The reconstructed shapes can then be processed with other data-driven methods for categorization, segmentation, matching and so on. Developing methods that can handle any 3D data representation, as well as jointly reconstructing and analysing shapes is a potential direction for future research we discuss in Section 11.

When polygon meshes are used as the input representation, an important aspect to consider is whether and how data-driven methods will deal with possible ‘defects’, such as non-manifold and non-orientable sets of polygons, inverted faces, isolated elements, self-intersections, holes and topological noise. The vast majority of meshes available in online repositories have these problems. Although there is a number of mesh repairing tools (see [CAK12] for a survey), they may not handle all different types of ‘defects’, and can take a significant amount of time to process each shape in a large data set. To avoid the issues caused by these ‘defects’, some data-driven methods uniformly sample the input meshes and work on the resulting point-based representation instead (e.g. [CKGK11, KLM\*13]).

**Data sets.** Although it is desirable to develop data-driven methods that can learn from a handful of training shapes or scenes, this is generally a challenging problem in machine learning [FFFP06]. Several data-driven methods in computer vision have been particularly successful due to the use of very large data sets that can reach the size of several millions of images [TFF08]. In contrast, data-driven approaches for 3D shape and scene processing approaches have mostly relied on data sets that reach the order of a few thousands so far (e.g. Princeton Shape Benchmark [SMKF04], or data sets collected from the web [KLM\*13]). Online repositories contain large amount of shapes, which can lead to the development of methods that will leverage data sets that are orders of magnitudes larger than the ones currently used. One significant example is the recently available ShapeNet [SSY\*15], which provides a richly annotated, large-scale data set of 3D shapes. Similar to ImageNet, a well-known image database in the computer vision community, ShapeNet is organized based on the WordNet hierarchy. It has indexed about 3 million models, out of which 220 000 models are classified into 3135 WordNet synsets (a synset refers to a meaningful concept in WordNet).

Another possibility is to develop synthetic data sets. A notable example is the pose and part recognition algorithm used in Microsoft’s Kinect that relies on 500K synthesized shapes of human bodies in different poses [SFC\*11]. In general, large data sets are important to capture the enormous 3D shape and scene variability, and can significantly increase the predictive performance and usability of learning



**Figure 2:** The general pipeline of data-driven geometry processing contains four major stages: data collection and pre-processing, feature extraction (or feature learning), learning and inference. The inference supports many applications which would produce new shapes or scenes through reconstruction modelling or synthesis. These new data, typically possessing labels for shapes or parts, can be used to enrich the input data sets and enhance the learning tasks in future, forming a data-driven geometry processing loop.

methods. A more comprehensive summary of existing online data collections can be found on our wikipedia [XKHK14].

## 2.2. 3D data processing and feature representation

It is common to perform some additional processing on the input representations of shapes and scenes before executing the main learning step. The reason is that the input representations of 3D shapes and scenes can have different resolutions (e.g. number of points or faces), scale, orientation and structure. In other words, the input shapes and scenes do not initially have any type of common parameterization or alignment. This is significantly different from other domains, such as natural language processing or vision, where text or image data sets frequently come with a common parameterization beforehand (e.g. images with the same number of pixels and objects of consistent orientation).

To achieve a common parameterization of the input shapes and scenes, one popular approach is to embed them in a *common geometric feature space*. For this purpose, a variety of shape descriptors have been developed. These descriptors can be classified into two main categories: *global shape descriptors* that convert each shape to a feature vector, and *local shape descriptors* that convert each point to a feature vector. Examples of global shape descriptors are Extended Gaussian Images [Hor84], 3D shape histograms [AKKS99, CK10], spherical functions [SV01], lightfield descriptors [CTSO03], shape distributions [OFCD02], symmetry descriptors [KFR04], spherical harmonics [KFR03], 3D Zernicke moments [NK03] and bags-of-words created out of local descriptors [BBOG11]. Local shape descriptors include surface curvature, Principal Component Analysis (PCA) descriptors, local shape diameter [SSCO08], shape contexts [BMP02, KHS10, KBLB12], spin images [JH99], geodesic distance features [ZMT05], heat-kernel descriptors [BBOG11] and depth features [SFC\*11]. Global shape descriptors are particularly useful for shape classification, retrieval and organization. Local shape descriptors are useful for partial shape matching, segmentation and point correspondence estimation. Before using any type of global or local descriptor, it is important to consider whether the descriptor will be invariant to different shape orientations, scales or poses. In the presence of noise and irregular mesh tessellations, it is important to robustly estimate local descriptors, since surface derivatives are particularly susceptible to surface

and sampling noise [KSNS07]. It is also possible to use several different descriptors as input, and let the learning step decide which ones are more relevant for each class of shapes [KHS10].

A different approach, which has attracted large attention in the computer vision community, is to avoid manually engineered features and instead directly learn features them from raw data. This approach has been enlightened by the recent developments in deep learning [Ben09, YN10], and in particular Convolutional Neural Networks (CNNs) [KSH12, SLJ\*15]. A number of deep learning architectures have been recently proposed to learn 3D shape and scene descriptors, operating on either voxel-based representations [WSK\*15], view-based projections [SMKLM15, XXS\*15], spectral representations [BMM\*15] or RGB-D data [SHB\*12, BSWR12, LBF13, BRF14].

Instead of embedding shapes in a common geometric feature space, several methods instead try to directly align shapes in Euclidean space. We refer the reader to the survey on dynamic geometry processing for a tutorial on rigid and non-rigid registration techniques [CLM\*12]. An interesting extension of these techniques is to include the alignment process in the learning step of data-driven methods, since it is interdependent with other shape analysis tasks such as shape segmentation and correspondences [KLM\*13].

Some data-driven methods require additional processing steps on the input. For example, learning deformation handles or fully generative models of shapes usually rely on segmenting the input shapes into parts with automatic algorithms [HKG11, SvKK\*11] and representing these parts with surface abstractions [YK12] or descriptors [KCKK12]. To decrease the amount of computation required during learning, it is also common to represent the shapes as a set of patches (super-faces) [HKG11] inspired by the computation of super-pixels in image segmentation.

## 2.3. Learning and inference

The processed representations of shapes and scenes are used to perform learning and inference for a variety of applications: shape classification, segmentation, matching, reconstruction, modelling, synthesis and scene analysis. The learning procedures significantly vary depending on the application, thus we discuss them individually in each of the following sections on these applications. As a common

theme, learning is viewed as an *optimization* problem that runs on a set of variables representing geometric, structural, semantic or functional properties of shapes and scenes. There is usually a single or multiple objective (or loss) functions for quantifying preferences for different models or patterns governing the 3D data. After learning a model from the training data, inference procedures are used to predict values of variables for new shapes or scenes. Again, the inference procedures vary depending on the application, and are discussed separately in the following sections. It is common that inference itself is an optimization problem, and sometimes is part of the learning process when there are latent variables or partially observed input shapes or scene data.

A general classification of the different types of algorithms used in data-driven approaches for shape and scene processing can be derived from the type of input information available during learning:

- **Supervised learning** algorithms are trained on a set of shapes or scenes annotated with labelled data. For example, in the case of shape classification, these labelled data can have the form of tags, while in the case of segmentation, the labelled data have the form of segmentation boundaries or part labels. The labelled data can be provided by humans or generated synthetically. After learning, the learned models are applied on different sets of shapes (test shapes) to produce results relevant to the task.
- **Unsupervised** algorithms co-analyse the input shapes or scenes without any additional labelled data, i.e. the desired output is unknown beforehand. The goal of these methods is to discover correlations in the geometry and structure of the input shape or scene data. For example, unsupervised shape segmentation methods usually perform some type of clustering in the feature space of points or patches belonging to the input shapes.
- **Semi-supervised** algorithms make use of shapes (or scenes) with and without any labelled data. Active learning is a special case of semi-supervised learning in which a learning algorithm interactively queries the user to obtain desired outputs for more data points related to shapes.

In general, supervised methods tend to output results that are closer to what a human would expect given the provided labelled data. However, they may fail to produce desirable results when the training shapes (or scenes) are geometrically and structurally dissimilar from the test shapes (or scenes). They also tend to require a substantial amount of labelled information as input, which can become a significant burden for the user. Unsupervised methods can deal with collections of shapes and scenes with larger variability and require no human supervision. However, they sometimes require parameter tuning to yield the desired results. Semi-supervised methods represent a trade-off between supervised and unsupervised methods: they provide more direct control to the user about the desired result compared to unsupervised methods, and often produce considerable improvements in the results by making use of both labelled and unlabelled shapes or scenes compared to supervised methods.

**The data-driven loop.** An advantageous feature of data-driven shape processing is that the output data, produced by learning and inference, typically come with rich semantic information. For example, data-driven shape segmentation produces parts with semantic

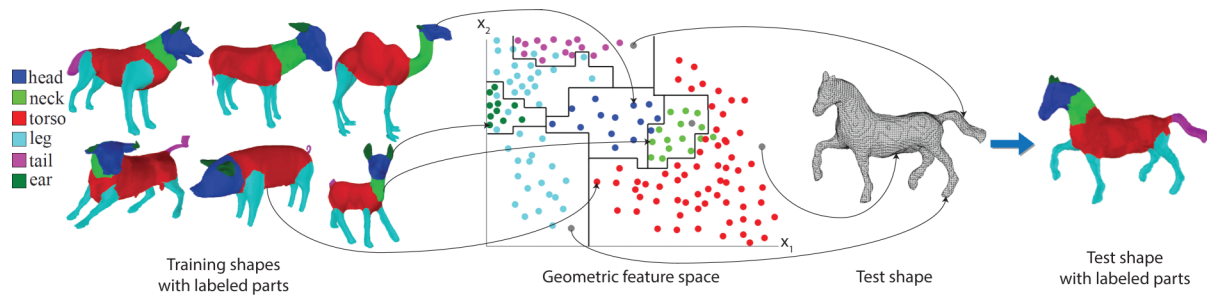
labels [KHS10]; data-driven reconstruction is commonly coupled with semantic part or shape recognition [SFCH12, NXS12]; data-driven shape modelling can generate readily usable shapes inheriting the semantic information from the input data [XZZ\*11]. These processed and generated data can be used to enrich the existing shape collections with both training labels and reusable contents, which in turn benefit subsequent learning. In a sense, data-driven methods *close the loop of data generation and data analysis* for 3D shapes and scenes; see Figure 2. Such concept has been practiced in several prior works, such as the data-driven shape reconstruction framework proposed in [PMG\*05] (Figure 12).

**Pipeline example.** To help the reader grasp the pipeline of data-driven methods, a schematic overview of the components is given in Figure 2. Depending on the particular application, the pipeline can have several variations, or some components might be skipped. We discuss the main components and steps of algorithms for each application in more detail in the following sections. A didactic example of the pipeline in the case of supervised shape segmentation is shown in Figure 3. The input shapes are annotated with labelled part information. A geometric descriptor is extracted for each point on the training shapes, and the points are embedded in a common feature space. The learning step uses a classification algorithm that non-linearly separates the input space into a set of regions corresponding to part labels in order to optimize classification performance (more details are provided in Section 4). Given a test shape, a probabilistic model is used to infer part labels for each point on that shape based on its geometric descriptor in the feature space.

#### 2.4. A comparative overview

Before reviewing the related works in detail, we provide a comparative overview of them in Table 5, and correlate them under a set of *criteria*:

- **Training data.** Data-driven methods can be categorized according to the shape or scene representations they operate on, the scale (size) of the training data sets they use, and the type of pre-processing applied to these data sets. The most common representation for shapes are polygon meshes and point clouds. 3D scenes are typically represented as an arrangement of individual shapes, usually organized in a scene graph. Pre-processing includes pre-segmentation, over-segmentation, pre-alignment, initial correspondence or/and labelling.
- **Features.** Roughly speaking, there are two types of feature representations involved in data-driven shape processing. The most commonly used feature representations are low-level ones, such as local geometric features (e.g. local curvature) and global shape descriptors (e.g. shape distribution [OFCD02]). If the input shapes are pre-segmented into meaningful parts, high-level structural representations (spatial relationships of parts) can be derived. Generally, working with high-level feature representations enables the learning of more powerful models for more advanced inference tasks, such as structural analysis [MWZ\*14], on complex man-made objects and scenes.
- **Learning model/approach.** The specific choice of learning method is often application-dependent. In most cases, machine learning techniques are adapted or developed from scratch to



**Figure 3:** Pipeline of a supervised segmentation algorithm [KHS10]. Given a set of shapes with labelled parts, the points of each shape are embedded in a common feature space based on their local geometric descriptors (a colour is assigned to points depending on their given part label). A classifier is learned to split the feature space into regions corresponding to each part label. Given a test shape, its points (shown in grey) are first embedded in the same space. Then part labels are inferred for all its points based on the learned classifier and an underlying structured probabilistic model (Section 4).

process geometric data. For some problems, such as shape correspondence, the core problem is to extract geometric correlations between different shapes in an unsupervised manner, which itself can be seen as a learning problem specific to geometry processing.

- **Learning type.** As discussed above, there are three basic types of data-driven methods, depending on the use of labelled training data: supervised, semi-supervised and unsupervised methods.
- **Learning outcome.** Learning can produce different types of outputs: parametric or non-parametric models (classifiers, clusterings, regressors, etc.), distance metrics which can be utilized for further analysis and/or feature representations learned from raw data.
- **Application.** The main applications of data-driven shape analysis and processing include classification, segmentation, correspondence, modelling, synthesis, reconstruction, exploration and organization.


### 3. Shape Classification

Data-driven techniques commonly make assumptions about the size and homogeneity of the input data set. In particular, existing analysis techniques often assume that all models belong to the same class of objects [KLM\*13] or scenes [FSH11], and cannot directly scale to entire repositories such as the Trimble 3D Warehouse [Tri14]. Similarly, techniques for data-driven reconstruction of indoor environments assume that the input data set only has furniture models [NXS12], while modelling and synthesis interfaces restrict the input data to particular object or scene classes [CKGK11, KCKK12, FRS\*12]. Thus, as a first step these methods need query a 3D model repository to retrieve a subset of relevant models.

Most public shape repositories such as 3D Warehouse [Tri14] rely on the users to provide tags and names of the shapes with little additional quality control measures. As a result, the shapes are sparsely labelled with inconsistent and noisy tags. This motivates developing automatic algorithms to infer text associated with models. Existing work focuses on establishing class memberships for an entire shape (e.g. this shape is a chair), as well as inferring finer scale attributes (e.g. this chair has a rocking leg).

**Classification** methods assign a class membership for unlabelled shapes. One approach is to retrieve for each unlabelled shape the most similar shape from a database of 3D models with known shape classes. There has been a large number of shape descriptors proposed in recent years that can be used in such a retrieval task, and one can refer to various surveys (e.g. [TV08]) for a thorough overview and comparisons. One can further improve classification results by leveraging machine learning techniques to learn classifiers that are based on global shape descriptors [FHK\*04, GKF09]. Barutcuoglu and DeCoro [BD06] demonstrate that Bayesian aggregation can be used to improve classification of shapes that are a part of a hierarchical ontology of objects. Geometry matching algorithms also facilitate distinguishing important features for classification [FS06, SF07]. Bronstein *et al.* [BBOG11] leverage ‘bag of features’ to learn powerful descriptor-space metrics for non-rigid shapes. These technique can be further improved by using sparse coding techniques [LBBC14]. In recent shape retrieval challenges, techniques based on bag of features demonstrated the best performance [LGA\*12, LLL\*14] in comparison to other alternatives. See Table 1 for a brief summary of some methods.

**Tag attributes** often capture fine-scale attributes of shapes that belong to the same class. These attributes can include presence or absence of particular parts, object style or comparative adjectives. Huang *et al.* [HSG13] developed a framework for propagating these attributes in a collection of partially annotated 3D models. For example, only brown models in Figure 4 were labelled, and blue models were annotated automatically. To achieve automatic labelling, they start by co-aligning all models to a canonical domain, and generate a voxel grid around the co-aligned models. For each voxel, they compute local shape features, such as spin images, for each shape. Then, they learn a distance metric that best discriminates between different tags. All shapes are finally embedded in a weighted feature space where nearest neighbours are connected in a graph. A graph cut clustering is used to assign tags to unlabelled shapes. Tag attributes can also be used to describe semantics, function or style of parts in shapes. Data-driven consistent segmentation and labelling techniques can be applied to propagate part tags across shapes (see Section 4). An alternative approach is to partition shapes into multiple sets of parts, then extract descriptors to define part



**Figure 4:** Fine-grained classification of 3D models [HSG13], where text labels are propagated from brown to blue models.


**Table 1:** Performance of several methods for shape classification (the accuracy in the right-most column as measured as fraction of correctly labelled shapes). Huang et al. [HSG13] predict fine-grained tag attributes for big collections of similar shapes. Golovinskiy et al. [GKF09] propose a method for classifying point clouds of objects in urban environments. The methods aimed at classifying meshes are evaluated on Princeton Shape Benchmark (PSB) [SF07, FS06, BD06]. To evaluate performance of the method in the presence of non-rigid deformations ShapeGoogle (SG) data set is also commonly used [BBOG11, LBBC14]. In addition, several methods can be found in regular shape retrieval challenges. The winner of 2012 challenge [LGA\*12] is [OF10] and the winner of 2014 challenge [LLL\*14] is DBSVC technique, both methods use bag of features for classification.

| Method   | Input data   | Shapes              | Classes | Acc |
|----------|--------------|---------------------|---------|-----|
| [HSG13]  | 3D Warehouse | 1206–5850           | 9–26    | 86  |
| [GKF09]  | LIDAR        | 1063                | 16      | 58  |
| [SF07]   | PSB          | 1814                | 90      | 75  |
| [FS06]   | PSB          | 1814                | 90      | 83  |
| [BD06]   | PSB          | 1814                | 90      | 84  |
| [BBOG11] | SG           | 715                 | 13      | 89  |
| [LBBC14] | SG           | 715                 | 13      | 91  |
| [LGA*12] | SHREC12      | 1200                | 60      | 88  |
| [LLL*14] | SHREC14      | 400–10 <sup>4</sup> | 19–352  | 87  |

similarity. A characteristic example of such an approach was demonstrated in Shapira et al. [SSS\*10]. Given the hierarchical segmentations of 3D shapes as input, part tagging was achieved by comparing local geometric features of parts as well as their context within the whole shape.

While the above method works well for discrete tags, they do not capture more continuous relations, such as animal A is more dangerous than animal B. Chaudhuri et al. [CKG\*13] focus on estimating ranking based on comparative adjectives. They use crowdsourcing to gather pairwise comparisons of shape parts with respect to different adjectives, and use a Support Vector Machine ranking method to predict attribute strengths from shape features for novel shape parts (Figure 5).

**Style similarity** methods have recently been proposed to classify shapes into style-related categories, e.g. buildings can be classified into architectural styles, such as Gothic, Baroque, Byzantine and so on. In contrast to the previously discussed approaches that rely on generic visual similarity measures to compare shapes, these methods



**Figure 5:** Ranking of parts with respect to ‘dangerous’ attribute (image from [CKG\*13]).

learn distance functions for style elements [LKS15] or common feature spaces [LHLF15] to quantify the stylistic similarity of shapes. The methods can be used to compare the style similarity of shapes, even when these belong to different classes (e.g. chairs and lamps). To align the style similarity measures with the human perception of style, style comparisons of shapes are gathered through crowdsourcing. The learned similarity measures can be used to retrieve stylistically similar shapes to populate a scene, or associate shapes with style-related tags.

While the techniques described above are suitable for retrieving and classifying shapes, a large number of applications require a more involved structural analysis to infer semantic and functional properties of shapes or their parts. The following two sections will discuss methods that perform structural analysis in collections of shapes based on segmentation and local matching.

#### 4. Shape Segmentation

The goal of data-driven shape segmentation is to partition the shapes of an input collection into parts, and also estimate part correspondences across these shapes. We organize the literature on shape segmentation into the following three categories: supervised segmentation, unsupervised segmentation and semi-supervised segmentation following the main classification discussed in Section 2. Table 2 summarizes representative techniques and reports their segmentation and part labelling performance based on established benchmarks. Table 3 reports characteristic running times for the same techniques.

**Table 2:** Performance of data-driven methods for segmentation in the Princeton Segmentation Benchmark (PSB) and COSEG data sets. Left to right: segmentation method, learning type depending on the nature of data required as input to the method, type of manual input if such required, segmentation performance expressed by the rand index metric [CGF09], labelling accuracy [KHS10] based on the PSB and COSEG data sets. We report the rand index segmentation error metric averaged over all classes of the PSB benchmark. The labelling accuracy is averaged over the Labelled PSB (L-PSB) benchmark excluding the ‘Bust’, ‘Mech’ and ‘Bearing’ classes. The reason is that there are no clear semantic correspondences between parts in these classes, or the ground-truth segmentations do not sufficiently capture semantic parts in their shapes. We report the labelling accuracy averaged over the categories of the COSEG data set used in [SvKK\*11]. The COSEG classes ‘iron’, ‘large chairs’, ‘large vases’, ‘tele-aliens’ were added later and are excluded here since most papers frequently do not report performance in those. We note that van Kaick et al. [vKTS\*11] reported the labelling accuracy in 10 of the L-PSB classes, while Wang et al. [WGW\*13] reported the labelling accuracy in seven of the L-PSB classes. The method by Kim et al. [KLM\*13] can run in either semi-supervised or unsupervised mode. In unsupervised mode, the corresponding labelling accuracy is 89.9% in the COSEG data set on average.

| Segmentation method | Learning type      | Type of manual input | PSB rand index (# train. shapes if applicable) | L-PSB accuracy (# train. shapes if applicable) | COSEG accuracy        |
|---------------------|--------------------|----------------------|--|--|-----------------------|
| [KHS10]             | Supervised         | Labelled shapes      | 9.4% (19)/14.8% (3)                            | 95.3% (19)/89.2% (3)                           | Unknown               |
| [BLVD11]            | Supervised         | Segmented shapes     | 8.8% (19)/9.7% (6)                             | Not applicable                                 | Not applicable        |
| [HKG11]             | Unsupervised       | None                 | 10.1%  | Not applicable                                 | Not applicable        |
| [SvKK*11]           | Unsupervised       | None                 | Unknown  | Unknown  | 87.7%                 |
| [vKTS*11]           | Supervised         | Labelled shapes      | Unknown  | ~88.7% (12), see caption                       | Unknown               |
| [HFL12]             | Unsupervised       | None                 | Unknown  | 88.5%  | 91.4%                 |
| [LCHB12]            | Semi-supervised    | Labelled shapes      | Unknown  | 92.3% (3)                                      | unknown               |
| [WAvK*12]           | Semi-supervised    | Link constraints     | Unknown  | Unknown  | ‘Close to error-free’ |
| [WGW*13]            | Supervised         | Labelled images      | Unknown  | ~88.0% (19), see caption                       | Unknown               |
| [KLM*13]            | Semi-/unsupervised | Box templates        | Unknown  | Unknown  | 92.7% (semi-superv.)  |
| [HWG14]             | Unsupervised       | None                 | Unknown  | Unknown  | 90.1%                 |
| [XSX*14]            | Supervised         | Labelled shapes      | 10.0%  | 86.0%  | Unknown               |
| [XXLX14]            | Supervised         | Labelled shapes      | 10.2% (19)                                     | 94.2 (19)/88.6 (5)                             | Unknown               |

**Table 3:** Running times reported for the data-driven segmentation methods of Table 2. We note that running times are reported in different data set sizes and processors in the referenced papers, while it is frequently not specified whether the execution uses one or multiple threads or whether the running times include all the algorithm steps, such as superface or feature extraction. Exact processor information is also frequently not provided. Thus, the reported running times of this table are only indicative and should not serve as a basis for a fair comparison.

| Segmentation method | Reported running times          | Data set size for reported running times | Reported processor       |
|---------------------|---------------------------------|--|--------------------------|
| [KHS10]             | 8 h train./5 min test.          | 6 train. shapes/1 test shape             | Intel Xeon E5355 2.66GHz |
| [BLVD11]            | 10 min train./1 min test.       | unknown for train./1 test shape          | Intel Core 2 Duo 2.99GHz |
| [HKG11]             | 32 h                            | 380 shapes                               | unknown, 2.4 GHz         |
| [SvKK*11]           | 10 min                          | 30 shapes                                | AMD Opteron 2.4GHz       |
| [vKTS*11]           | 10 h train./few min test.       | 20-30 train. shapes/1 test shape         | AMD Opteron 1GHz         |
| [HFL12]             | 8 min (excl. feat. extr.)       | 20 shapes                                | Intel dual-core 2.93GHz  |
| [LCHB12]            | 7 h train./few min test.        | 20 shapes                                | Intel I7 2600 3.4GHz     |
| [WAvK*12]           | 7 min user interaction          | 28 shapes                                | unknown                  |
| [WGW*13]            | 1.5 min (no train. step)        | 1 test shape                             | unknown                  |
| [KLM*13]            | 11 h                            | 7442 shapes                              | unknown                  |
| [HWG14]             | 33 h                            | 8401 shapes                              | unknown, 3.2GHZ          |
| [XSX*14]            | 30 s (no train. step)           | 1 test shape                             | Intel I5 CPU             |
| [XXLX14]            | 15 s train. (excl. feat. extr.) | 6 train. shapes                          | Intel Quad-Core 3.2 GHz  |

#### 4.1. Supervised shape segmentation


**Classification techniques.** Supervised shape segmentation is frequently formulated as a classification problem. Given a training set of shapes containing points, faces or patches that are labelled according to a part category (see Figure 3), the goal of a classifier is to identify which part category other points, faces or patches from different shapes belong to. Supervised shape segmentation is executed in two steps: during the first step, the parameters of the classifier are

learned from the training data. During the second step, the classifier is applied on new shapes. A simple linear classifier has the form:

$$c = f\left(\sum_j \theta_j \cdot x_j\right), \quad (1)$$

where  $x_j$  is a geometric feature of a point (face, or patch), such as the ones discussed in Section 2. The parameters  $\theta_j$  serve as





**Figure 6:** A random forest classifier applied on depth data representing a human body shape (image from [FGG\*13]).

weights for each geometric feature. The function  $f$  is non-linear and maps to a discrete value (label), which is a part category, or to probabilities per category. In general, choosing a good set of geometric features that help predicting part labels, and employing classifiers that can discriminate the input data points correctly are important design choices. There is no rule of thumb on which is the best classifier for a problem. This depends on the underlying distribution and characteristics of the input geometric features, their dimensionality, amount of labelled data, existence of noise in the labelled data or shapes, training and test time constraints—for a related discussion on how to choose a classifier for a problem, we refer the reader to [MRS08]. Due to the large dimensionality and complexity of geometric feature spaces, non-linear classifiers are more commonly used. For example, to segment human bodies into parts and recognize poses, the Microsoft’s Kinect uses a random forest classifier trained on synthetic depth images of humans of many shapes and sizes in highly varied poses sampled from a large motion capture database [SFC\*11] (Figure 6).

**Structured models.** For computer graphics applications, it is important to segment shapes with accurate and smooth boundaries. For example, to help the user create a new shape by recombining parts from other shapes [FKS\*04], irregular and noisy segmentation boundaries can cause problems in the part attachment. From this aspect, using a classifier per point/face independently is usually not enough. Thus, it is more common to formulate the shape segmentation problem as an energy minimization problem that involves a unary term assessing the consistency of each point/face with each part label, as well as a pairwise term assessing the consistency of neighbouring points/faces with pairs of labels. For example, pairs of points that have low curvature (i.e. are on flat surface) are more likely to have the same part label. This energy minimization formulation has been used in several single-shape and data-driven segmentations (unsupervised or supervised) [KT03, ATC\*05, SSS\*10, KHS10]. In the case of supervised segmentation [KHS10], the energy can be written as

$$E(\mathbf{c}; \theta) = \sum_i E_{\text{unary}}(c_i; \mathbf{x}_i, \theta_1) + \sum_{i,j} E_{\text{pairwise}}(c_i, c_j; \mathbf{y}_{ij}, \theta_2), \quad (2)$$

where  $\mathbf{c} = \{c_i\}$  is a vector of random variables representing the part label per point (or face)  $i$ ,  $\mathbf{x}_i$  is its geometric feature vector,  $i, j$  are indices to points (or faces) that are considered neighbours,  $\mathbf{y}_{ij}$  is a geometric feature vector representing dihedral angle, angle between normals or other features and  $\theta = \{\theta_1, \theta_2\}$  are the energy parameters. The important difference of supervised data-driven methods with previous single-shape segmentation methods is that the parameters  $\theta$  are automatically learned from the training shapes to capture

complex feature space patterns per part [ATC\*05, KHS10]. We also note that the above energy of Equation (2), when written in an exponentiated form and normalized, can be treated as a probabilistic graphical model [KF09], called Conditional Random Field [LMP01] that represents the joint probability distribution over part labels conditioned on the input features:

$$P(\mathbf{c}|\mathbf{x}, \mathbf{y}, \theta) = \exp(-E(\mathbf{c}; \theta))/Z(\mathbf{x}, \mathbf{y}, \theta), \quad (3)$$

where  $Z(\mathbf{x}, \mathbf{y}, \theta)$  is a normalization factor, also known as partition function. Minimizing the energy of Equation (2), or correspondingly finding the assignment  $\mathbf{c}$  that maximizes the above probability distribution is known as a Maximum *a Posteriori* (MAP) inference problem that can be solved in various manners, such as graph cuts, belief propagation, variational or linear programming relaxation techniques [KF09].

The parameters  $\theta$  can be jointly learned through maximum likelihood (ML) or MAP estimates [KF09]. However, due to high computational complexity of ML or MAP learning and the non-linearity of classifiers used in shape segmentation, it is common to train the parameters  $\theta_1$  and  $\theta_2$  of the model separately, i.e. train the classifiers of the unary and pairwise term separately [SM05]. The exact form of the unary and pairwise terms vary across supervised shape segmentation methods: the unary term can have the form of a log-linear model [ATC\*05], cascade of JointBoost classifiers [KHS10], Gentleboost [vKTS\*11] or feedforward neural networks [XXLX14]. The pairwise term can have the form of a learned log-linear model [ATC\*05], label-dependent GentleBoost classifier [KHS10] or a smoothness term based on dihedral angles and edge length tuned by experimentation [SSS\*10, vKTS\*11, XXLX14]. Again the form of the unary and pairwise terms depend on the amount of training data, dimensionality and underlying distribution of geometric features used and computational cost.

**Joint labelling.** Instead of applying the learned probabilistic model to a single shape, an alternative approach is to find correspondences between faces of pairs of shapes, and incorporate a third ‘intershape’ term in the energy of Equation (2) [vKTS\*11]. The ‘intershape’ term favours pairs of corresponding faces on different shapes to have the same label. As a result, the energy can be minimized jointly over a set of shapes to take into account any additional correspondences.

**Boundary learning.** Instead of applying a classifier per mesh point, face or patch to predict a part label, a different approach is to predict the probability that a polygon mesh edge is a segmentation boundary [BLVD11]. The problem can be formulated as a binary classifier (e.g. Adaboost) that is trained from human segmentation boundaries. The input to the classifier are geometric features of edges, such as dihedral angles, curvature and shape diameter. The output is a probability for an edge to be a segmentation boundary. Since the predicted probabilities over the mesh do not correspond to closed smooth boundaries, a thinning and an active contour model [KWT88] are used in post-processing to produce the final segmentations.

**Transductive segmentation.** Another way to formulate the shape segmentation problem is to group patches on a mesh such that the

segment similarity is maximized between the resulting segments and the provided segments in the training database. The segment similarity can be measured as the reconstruction cost of the resulting segment from the training ones. The grouping of patches can be solved as an integer programming problem [XSX\*14].

**Shape segmentation from labelled images.** Instead of using labelled training shapes for supervised shape segmentation, an alternative source of training data can come in the form of segmented and labelled images, as demonstrated by Wang *et al.* [WGW\*13]. Given an input 3D shape, this method first renders 2D binary images of it from different viewpoints. Each binary image is used to retrieve multiple segmented and labelled training images from an input database based on a bi-class Hausdorff distance measure. Each retrieved image is used to perform label transfer to the 2D shape projections. All labelled projections are then backprojected onto the input 3D model to compute a labelling probability map. The energy function for segmentation is formulated by using this probability map in the unary term expressed per face or point, while dihedral angles and Euclidean distances are used in the pairwise term.

#### 4.2. Semi-supervised shape segmentation

**Entropy regularization.** The parameters  $\theta$  of Equation (2) can be learned not only from the training labelled shapes, but also from the unlabelled shapes [LCHB12]. The idea is that learning should maximize the likelihood function of the parameters over the labelled shapes, and also minimize the entropy (uncertainty) of the classifier over the unlabelled shapes (or correspondingly maximize the negative entropy). The idea is that minimizing the entropy over unlabelled shapes encourages the algorithm to find putative labellings for the unlabelled data [JWL\*06]. However, it is generally hard to strike a balance between the likelihood and entropy terms.

**Metric embedding and active learning.** A more general formulation for semi-supervised segmentation was presented in [WAVK\*12]. Starting from a set of shapes that are co-segmented in an unsupervised manner [SvKK\*11], the user interactively adds two types of constraints: ‘must-link’ constraints, which specify that two patches (super-faces) should belong to the same cluster, and ‘cannot-link’ constraints which specify that two patches must be in different clusters. These constraints are used to perform constrained clustering in an embedded feature space of super-faces coming from all the shapes of the input data set. The key idea is to transform the original feature space, such that super-faces with ‘must-link’ constraints come closer together to form a cluster in the embedded feature space, while super-faces with ‘cannot-link’ constraints move away from each other. To minimize the effort required from the user, the method suggests to the user pairs of points in feature space that when constrained are likely to improve the co-segmentation. The suggestions involve points that are far from their cluster centres, and have a low confidence of belonging to their clusters.

**Template fitting.** A different form of partial supervision can come in the form of part-based templates. Kim *et al.*'s method [KLM\*13] allows users to specify or refine a few templates made out of boxes representing expected parts in an input database. The boxes itera-

tively fit to the shapes of a collection through simultaneous alignment, surface segmentation and point-to-point correspondences estimated between each template and each input shape. Alternatively, the templates can be inferred automatically from the shapes of the input collection without human supervision based on single shape segmentation heuristics. Optionally, the user can refine and improve these estimated templates. From this aspect, Kim *et al.*'s method can run in either a semi-supervised or unsupervised method. It was also the first method to handle segmentation and correspondences in collections with size on the order of thousands of shapes.


#### 4.3. Unsupervised segmentation

Unsupervised data-driven shape segmentation techniques fall into two categories: clustering-based techniques and matching-based techniques. In the following, we highlight the key idea of each type of approach.

**Clustering-based techniques** are adapted from supervised techniques. They compute feature descriptors on points or faces. Clustering is performed over all points/faces over all shapes. Each resulting cluster indicates a consistent segment across the input shapes. The promise of the clustering-based approach is that when the number of shapes becomes large, the sampling density in the clustering space becomes dense enough, so that certain statistical assumptions are satisfied, e.g. diffusion distances between points from different clusters is significantly larger than those between points within each cluster. When these assumptions are satisfied, clustering-based approaches may produce results that are comparable to supervised techniques (*cf.* [HFL12]). In [SvKK\*11], the authors utilize spectral clustering to perform clustering. In [HFL12], the authors employ subspace clustering, a more advanced clustering method, to obtain improved results.

Clustering methods can also be applied to shape parts. In [XLZ\*10], the authors perform co-analysis over a set of shapes via factoring out the part scale variation by grouping the shapes into different styles, where style is defined by the anisotropic part scales of the shapes. In [vKXZ\*13], the authors introduce unsupervised co-hierarchical analysis of a set of shapes. They propose a novel cluster-and-select scheme for selecting representative part hierarchies for all shapes and grouping the shapes according to the hierarchies. The method can be used to compute consistent hierarchical segmentations for the input set.

**Matching-based methods** [GF09, HKG11, WHG13, HWG14] build maps across shapes and utilize these maps to achieve consistency of segmentations. As shown in Figure 7, this strategy allows us to identify meaningful parts despite the lack of strong geometric cues on a particular shape. Likewise, the approach is able to identify coherent single parts even when the geometry of the individual shape suggests the presence of multiple segments. A challenge here is to find a suitable shape representation so that maps across diverse shapes are well-defined. In [HKG11], Huang *et al.* introduce an optimization strategy that jointly optimizes shape segmentations and maps between optimized segmentations. Since the maps are defined at the part level, this technique is suitable for heterogeneous shape collections. Experimentally, it generates comparable results with supervised method [KHS10] on the Princeton segmentation benchmark. Recently, Huang *et al.* [HWG14] formulated the same



**Figure 7:** Comparison of single-shape segmentation (left) and joint shape segmentation (right) on models from the PSB benchmark [CGF09]. Each segmentation on the left was produced by the top-performing algorithm in the benchmark for that shape. The segmentations on the right were produced by [HKG11], which jointly optimized segmentations and correspondences across the entire data set.

idea under the framework of functional maps [OBSC\*12] and gain improved segmentation quality and computational efficiency.

## 5. Joint Shape Matching

Another fundamental problem in shape analysis is shape matching, which finds relations or maps between shapes. These maps allow us to transfer information across shapes and aggregate information from a collection of shapes for a better understanding of individual shapes (e.g. detecting shared structures such as skeletons or shape parts). They also provide a powerful platform for comparing shapes (i.e. with respect to different measures and at different places). As we can see from other sections, shape maps are widely applied in shape classification and shape exploration as well.

So far, most existing research in shape matching has focused on matching pairs of shapes in isolation. We refer to [vKZHC011] for a survey and to [LH05, LF09, vKZHC011, OMMG10, KLF11, OBSC\*12] for recent advances. Although significant progress has been made, state-of-the-art techniques are limited to shapes that are similar to each other. On the other hand, these techniques tend to be insufficient among shape collections that possess large geometric and topological variations.

The availability of large shape collections offers opportunities to address this issue. Intuitively, when matching two dissimilar shapes, we may utilize intermediate shapes to transfer maps. In other words, we can build maps between similar shapes, and use the composite maps to obtain maps between less similar shapes. As we will see shortly, this intuition can be generalized to enforcing a cycle-consistency constraint; namely, composite maps along cycles should be the identity map and the composite map between any two shapes is path-independent. In this section, we discuss joint shape matching techniques that take a shape collection and

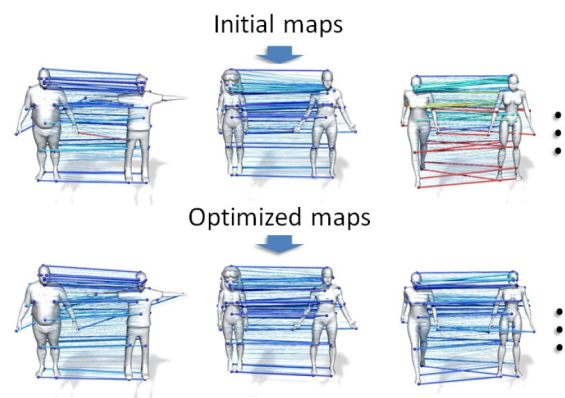
noisy maps as input, and output improved maps across the shape collection.

### 5.1. Model graph and cycle consistency

To formulate the joint matching problem, we consider a model graph  $\mathcal{G} = (\mathcal{S}, \mathcal{E})$  (cf. [Hub02]). The vertex set  $\mathcal{S} = \{S_1, \dots, S_n\}$  consists of the input shapes. The edge set  $\mathcal{E}$  characterizes the pairs of shapes that are selected for performing pairwise matching. For small-scale data sets, we typically match all pairs of shapes. For large-scale data sets, the edge set usually connects shapes that are similar according to a pre-defined shape descriptor [KLM\*12, HSG13], thus generating a sparse shape graph.

The key component of a joint matching algorithm is to utilize the so-called cycle-consistency constraint. In particular, if all the maps in  $\mathcal{G}$  are correct, then composite maps along any loops should be the identity map. This is true for maps that are represented as transformations (e.g. rotations and rigid/affine transformations), or full pointwise maps that can be described as permutation matrices). We can easily modify the constraint to handle partial maps; namely, each point, when transformed along a loop, either disappears or goes back to the original point (see Figure 8 and [HWG14] for details).

The cycle-consistency constraint is useful because the initial maps, which are computed between pairs of shapes in isolation, are not expected to satisfy the cycle-consistency constraint. On the other hand, although we do not know which maps or correspondences are incorrect, we can detect inconsistent cycles. These inconsistent cycles provide useful information for us to detect incorrect correspondences or maps, i.e. an inconsistent cycle indicates that at least one of the participating maps or correspondences is incorrect. To incorporate this observation into algorithms, one has to formulate the cycle-consistency constraint properly. Existing works in data-driven shape matching fall into two categories: combinatorial techniques and matrix recovery-based techniques. The remainder of this section provides the details.



**Figure 8:** Joint shape matching takes as input maps computed between pairs of shapes in isolation and utilizes the cycle-consistency constraint to improve shape maps. This figure shows the result of Huang et al. [HWG14], which performs joint shape matching under the functional map setting.

## 5.2. Combinatorial techniques

**Spanning tree optimization.** Earlier works in joint matching aim at finding a spanning tree in the model graph. In [GMB04, HFG\*06], the authors propose to use the maximum spanning tree (MST) of the model graph. However, this strategy can easily fail since a single incorrect edge in the MST may break the entire matching result. In the seminal work [Hub02], Huber showed that finding the best spanning tree maximizing the number of consistent edges is NP-hard. Although finding the best spanning tree is not tractable, Huber introduced several local operations for improving the score of spanning trees. However, these approaches are generally limited to small-scale problems so that the search space can be sufficiently explored.

**Inconsistent cycle detection.** Another line of approaches [ZKP10, RSSL11, NBCW\*11] applies global optimization to select cycle-consistent maps. These approaches are typically formulated as solving constrained optimization problems, where objective functions encode the scores of selected maps, and constraints enforce the consistency of selected maps along cycles. The major advantage of these approaches is that the correct maps are determined globally. However, as the cycle-consistency constraint needs to apportion blame along many edges on a cycle, the success of these approaches relies on the assumption that correct maps are dominant in the model graph so that the small number of bad maps can be identified through their participation in many bad cycles.

**Markov Random Field (MRF) formulation.** Joint matching may also be formulated as solving a second-order Markov Random Field (or MRF) [CAF10b, CAF10a, COSH11, HZG\*12]. The basic idea is to sample the transformation/deformation space of each shape to obtain a candidate set of transformation/deformation samples per shape. Joint matching is then formulated as optimizing the best sample for each shape. The objective function considers initial maps. Specifically, each pair of samples from two different shapes would generate a candidate map between them. The objective function then formulates second-order potentials, where each term characterizes the alignment score between these candidate maps and the initial maps [HSG13, HZG\*12].

The key challenge in the MRF formulation is generating the candidate samples for each shape. The most popular strategy is to perform uniform sampling [COSH11, HSG13], which works well when the transformation space is low-dimensional. To apply the MRF formulation on high-dimensional problems, Huang *et al.* [HZG\*12] introduce a diffusion-and-sharpening strategy. The idea is to diffuse the maps among the model graph to obtain rich samples of candidate transformations or correspondences and then perform clustering to reduce the number of candidate samples.

## 5.3. Matrix-based techniques

A recent trend in map computation is to formulate joint map computation as inferring matrices [SW11, HZG\*12, KLM\*12, HG13, WS13, CGH14, HWG14]. The basic idea is to consider a big map collection matrix

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} & \cdots & \mathbf{X}_{1n} \\ \mathbf{X}_{21} & \mathbf{X}_{22} & \cdots & \mathbf{X}_{2n} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{X}_{n1} & \cdots & \cdots & \mathbf{X}_{nn} \end{bmatrix},$$

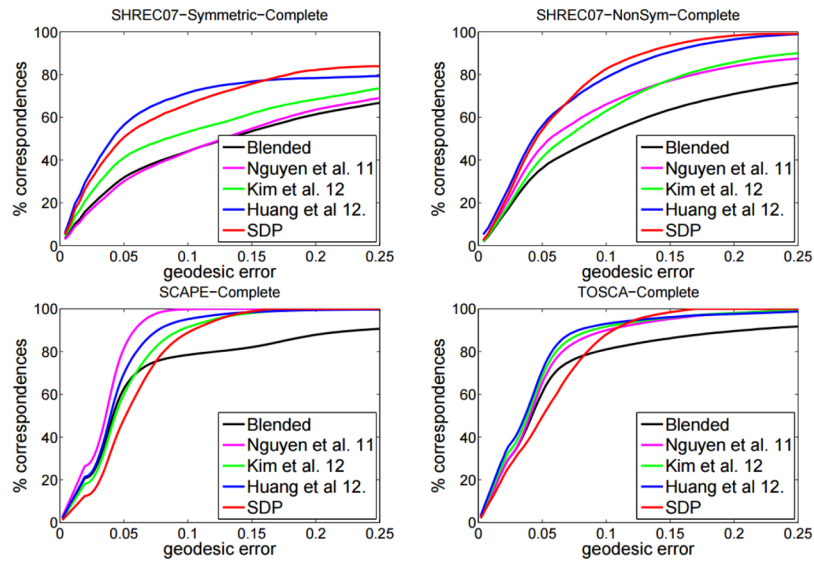
where each block  $\mathbf{X}_{ij}$  encodes the map from shape  $S_i$  to shape  $S_j$ . In this matrix representation, the cycle-consistency constraint can be equivalently described by simple properties of  $\mathbf{X}$ , i.e. depending on the types of maps,  $\mathbf{X}$  is either positive semidefinite or low rank (*cf.* [HG13, HWG14]). In addition, we may view the initial pairwise maps as noisy measurements of the entries of  $\mathbf{X}$ . Based on this perspective, we can formulate joint matching as matrix recovery from noisy measurements of its entries.

**Spectral techniques.** The initial attempts in matrix recovery are spectral techniques and their variants [SW11, KLM\*12, WHG13]. The basic idea is to consider the map collection  $\mathbf{X}^{\text{input}}$  that encodes initial maps in its blocks. Then, the recovered matrix is given by  $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$ , where  $\mathbf{U}$ ,  $\Sigma$ ,  $\mathbf{V}$  represent the singular value decomposition (or SVD) of  $\mathbf{X}^{\text{input}}$ . Various methods have added heuristics on top of this basic procedure. For example, Kim *et al.* [KLM\*12] use the optimized maps to recompute initial maps.

This SVD strategy can be viewed as matrix recovery because  $\mathbf{X}$  is equivalent to the optimal low-rank approximation of  $\mathbf{X}^{\text{input}}$  (with given rank) under the matrix Frobenius norm. However, as the input maps may contain outliers, employing the Frobenius norm for matrix recovery is suboptimal. Moreover, it is hard to analyse these techniques, even in the very basic setting where maps are given by permutation matrices [PKS13].

**Point-based maps.** In a series of works, Huang and coworkers [HG13, CGH14, HCG14] consider the case of point-based maps and develop joint matching algorithms that admit theoretical guarantees. The work of [HG13] considers the basic setting of permutation matrix maps and proves the equivalence between cycle-consistent maps and the low-rank or positive semi-definiteness of the map collection matrix. This leads to a semidefinite programming formulation for joint matching. In particular, the L1 norm is used to measure the distance between the recovered maps and the initial maps. The authors provide exact recovery conditions, which state that the ground-truth maps can be recovered if the percentage of incorrect correspondences in the input maps is below a constant. In a follow-up work, Chen *et al.* [CGH14] extends this to partial maps and provide a better analysis in the case where incorrect correspondences in the input maps are random. The computational issue is addressed in [HCG14], which employs the alternating direction of multiplier methods for optimization. Figure 9 compares the SDP formulation of [HG13] with several other data-driven shape matching methods. Note that all data-driven shape matching methods improve upon pairwise matching, yet the SDP formulation leads to the largest improvement.

**Rotations and functional maps.** Maps that are represented by general matrices (e.g. rotations or functional maps) can also be handled in a similar fashion. In [WS13], Wang and Singer consider



**Figure 9:** Comparison among various data-driven shape matching methods: optimized composite maps [NBCW\*11], fuzzy correspondences [KLM\*12], hub-and-spoke network [HZG\*12] and semi-definite programming relaxation [HG13]. The input maps are given by blended intrinsic maps [KLF11].

the case of rotations between objects. Their formulation is similar to [HG13] but utilize an L1 Frobenius norm for measuring the distance between initial rotations and recovered rotations. Recently, Huang *et al.* [HWG14] extend this idea to functional maps. The major difference between functional maps and point-based maps or rotations is that the map collection matrix is no-longer symmetric. Thus, their method is formulated to recover low-rank matrices.

#### 5.4. Discussion and future directions

The key to joint shape matching is to have a proper formulation of the cycle-consistency constraint. We have witnessed the evolution of this constraint from earlier works on combinatorial search and inconsistent cycle detection to more recent works which use spectral techniques, MRF-based methods and matrix recovery techniques. In particular, matrix recovery techniques admit theoretical guarantees, providing a fundamental understanding of why joint shape matching can improve upon isolated pairwise matching.

One future direction is to integrate pairwise matching and joint matching into one optimization problem. Since the major role of joint matching is to remove the noise presented in pairwise matching, it makes sense to perform them together. Such unified approaches have the potential to further improve upon decomposed approaches (i.e. from pairwise to joint). The technical challenge is to find map representations so that pairwise matching and map consistency can be formulated in the same framework.

## 6. Shape Reconstruction

Reconstructing geometric shapes from physical objects is a fundamental problem in geometry processing. The input to this problem is usually a point cloud produced by aligned range scans, which provides an observation of an object. The goal of a shape reconstruction algorithm is to convert this point cloud into a high-quality


geometric model. In practice, the input point cloud data are noisy and incomplete. Thus, the key to a successful shape reconstruction algorithm is formulating appropriate shape priors. Traditional shape reconstruction algorithms usually utilize generic priors, such as surface smoothness [DTB06], and typically assume that the input data capture most of the object’s surface. To handle higher degrees of noise and partiality of the input data, it is important to build structural shape priors.

Data-driven techniques tackle this challenge by leveraging shape collections to learn strong structural priors from similar objects, and use them to reconstruct high-quality 3D models. Existing approaches fall into two categories, based on how they represent the shape priors: *parametric* and *non-parametric*. The former usually builds a low-dimensional parametric representation of the underlying shape space, learning the representation from exemplars and enforcing the parameterization when reconstructing new models. Parametric methods typically require building correspondences across the exemplar shapes. In contrast, non-parametric methods directly operate on the input shapes by copying and deforming existing shapes or shape parts.

### 6.1. Parametric methods

**Morphable face.** The morphable face model [BV99] is a representative work of parametric data-driven shape reconstruction, a technique which reconstructs 3D textured faces from photos and scans. The model is learned from a data set of prototypical 3D shapes of faces, and can then be used to derive a new 3D face from a novel image (see Figure10).

In particular, the morphable face model represents the geometry of a face with a shape-vector  $S = (\mathbf{p}_1^T, \dots, \mathbf{p}_n^T)^T \in \mathbb{R}^{3n}$ , which contains the 3D coordinates of its  $n$  vertices. Similarly, it encodes the texture of a face by a texture-vector  $T = (\mathbf{c}_1^T, \mathbf{c}_2^T, \dots, \mathbf{c}_n^T) \in \mathbb{R}^{3n}$ ,



**Figure 10:** Derived from a data set of prototypical 3D scans of faces, the morphable face model contributes to two main steps in face manipulation: (1) deriving a 3D face model from a novel image, and (2) modifying shape and texture in a natural way [BV99].

that contains the RGB colour values of the corresponding vertices. A morphable face model is then constructed using a database of  $m$  exemplar faces, each represented by its shape-vector  $S_i$  and  $T_i$ . In [BV99], the exemplar faces are constructed by matching a template to scanned human faces.

The morphable face model uses PCA to characterize the shape space. A new shape and its associated texture are given by

$$S_{\text{mod}} = \bar{S} + \sum_{i=1}^{m-1} \alpha_i s_i, \quad T_{\text{mod}} = \bar{T} + \sum_{i=1}^{m-1} \beta_i t_i,$$


where  $\bar{S}$  and  $\bar{T}$  are the mean-shape and mean-texture, respectively, and  $s_i$  and  $t_i$  are eigenvectors of the covariance matrices.  $\alpha_i$  and  $\beta_i$  are coefficients. PCA also gives probability distributions over coefficients. The probability for coefficient  $\alpha_i$  is given by

$$p(\{\alpha_i\}) \sim \exp\left(-\frac{1}{2} \sum_{i=1}^{m-1} (\alpha_i / \sigma_i)^2\right)$$

with  $\sigma_i^2$  being the corresponding eigenvalue of the shape covariant matrix  $C_S$  (the probability  $p(\{\beta_i\})$  is computed in a similar way).

With this morphable face model, reconstruction of textured models can be posed as a small-scale non-linear optimization problem. For example, given a 2D image of a human face  $I_{\text{input}}$ , one can reconstruct the underlying textured 3D model by searching for a similar rendered face  $I(\{\alpha_i\}, \{\beta_i\}, p)$ , parameterized by the shape and texture coefficients  $\alpha_i$  and  $\beta_i$ , and the rendering parameters  $p$  (e.g. camera configuration, lighting parameters). The optimization problem is formulated as minimizing a data term, which measures the distance between the input image and the rendered image, and regularization terms that are learned from exemplar faces. The success of the morphable model relies on the low-dimensionality of the solution space, and so is also applicable to several other data sets where this assumption holds, including the domain of human bodies and poses.

**Morphable human bodies.** Allen *et al.* [ACP03] generalize the morphable model to characterize human bodies (Figure 11). Given a set of 250 scanned human bodies, the method first performs non-rigid registration to fit a hole-free, artist-generated mesh (template)



**Figure 11:** Parameterizing the variation in human shapes can be used to synthesize new individuals or edit existing ones [ACP03].

to each of these scans. The result is a set of mutually consistent parameterized shapes based on the corresponding vertex positions originating from the template. Similar to [BV99], the method employs PCA to characterize the shape space, which enables applications in shape exploration, synthesis and reconstruction.

In addition to variations in body shape, human models also exhibit variations in poses. The SCAPE model (Shape Completion and Animation for People) [ASK\*05] addresses this challenge by learning two separate models of body deformation—one accounting for variations in poses and one for differences in body shapes. The pose deformation component is acquired from a set of dense 3D scans of a single person in multiple poses. A key aspect of the pose model is that it decomposes deformation into a rigid and a non-rigid component. The rigid component is modelled using a standard skeleton system. The non-rigid component, which captures remaining deformations such as flexing of the muscles, associates each triangle with a local affine transformation matrix. These transformation matrices are learned from exemplars using a joint regression model. In [HSS\*09], Hasler *et al.* introduce a unified model for parameterizing both shapes and poses. The basic idea is to consider the relative transformations between all pairs of neighbouring triangles. These transformation matrices allow us to reconstruct the original shape by solving a least-squares problem. In this regard, each shape is encoded as a set of edgewise transformation matrices, which are fit into the PCA framework to obtain a statistical model of human shapes. The model is further extended to estimate shapes of dressed humans from range scans [HSR\*09].

Recent works on statistical human shape analysis focus on combining learned shape priors with sparse observations and special effects. In [LMB14], the authors introduce an approach that reconstructs high-quality shapes and poses from a sparse set of markers. The success of this approach relies on learning meaningful shape priors from a database consisting of thousands of shapes. In [TMB14], the authors study human breathing from acquired data.

**Data-driven tracking.** Object tracking aims to analyse dynamic shapes and/or poses of physical objects. Successful tracking techniques (e.g. [WLVGP09, WBLP11, LYYB13, CWLZ13, CHZ14]) typically utilize parametric shape spaces. These reduced shape spaces provide shape priors that improve both the efficiency and robustness of the tracking process. The way to utilize and construct shape spaces vary in different settings, and are typically tailored to the specific problem setting. Weise *et al.* [WLVGP09] utilize a linear PCA subspace trained with a very large set of pre-processed facial expressions. This method requires an extended training session with a careful choice of facial action units. In addition, the learned face model is actor-specific. These restrictions are partially

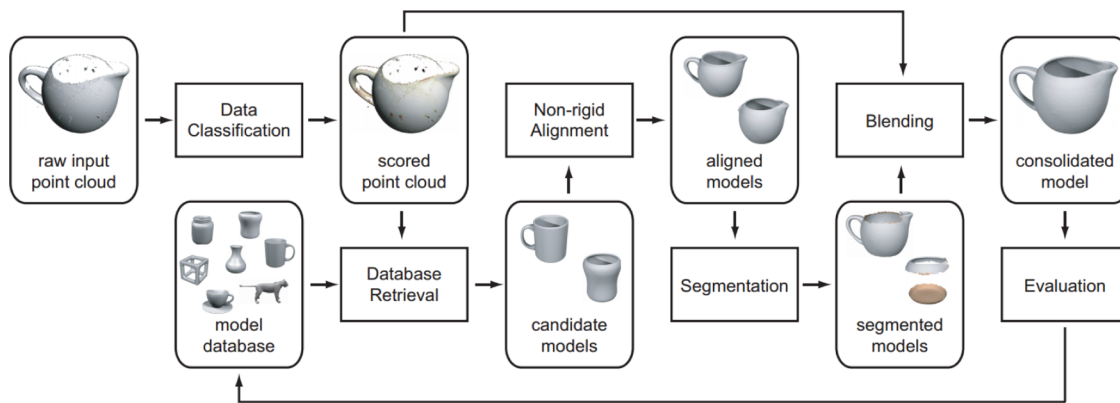


Figure 12: The data-driven shape reconstruction pipeline proposed in [PMG\*05].

resolved in [LWP10], which introduces an example-based blendshape optimization technique, involving only a limited number of random facial expressions. In [WBLP11], the authors combine both blendshapes and data-driven animation priors to improve the tracking performance. In a recent work, Li *et al.* [LYYB13] employ adaptive PCA to further improve tracking performance on nuanced emotions and micro-expression. The key idea is to combine a general blendshape PCA model and a corrective PCA model that is updated on-the-fly. This corrective PCA model captures the details of the specific actor and missing deformations from the initial blendshape model.

## 6.2. Non-parametric methods

Parametric methods require canonical domains to characterize the shape space, which have so far been demonstrated within domains of organic shapes, such as body shapes or faces. In this section, we discuss another category of methods that have shown the potential to handle more diverse shape collections.

Generally speaking, a non-parametric data-driven shape reconstruction method utilizes a collection of relevant shapes and combines three phases, i.e. a query phase, a transformation phase and an assembly phase. Existing methods differ in how the input shape collection is pre-processed and how these phases are performed.

**Example-based scan completion.** Pauly *et al.* [PMG\*05] introduce one of the first non-parametric systems. As shown in [PMG\*05], the method takes a point cloud and a collection of complete objects as input. The reconstruction procedure incorporates all three phases described above. The first phase determines a set of similar objects. The retrieval phase combines both text-based search and PCA signatures, which are then refined by rigid alignment. The second step performs non-rigid alignment between the retrieved shapes and the input point cloud. This step partitions the input point cloud into a set of patches, where each patch is associated with one retrieved shape (via the corresponding region). The final phase merges the corresponding regions into a unified shape.

Nan *et al.* [NXS12] introduce a similar system for indoor scene reconstruction. Given an input point cloud of an indoor scene that con-

sists of a set of objects with known categories, the method searches in a database of 3D models to find matched objects and then deforms them in a non-rigid manner to fit the input point cloud. Note that this method treats complete 3D objects as building blocks, so the final reconstruction does not necessarily reflect the original scene. Shao *et al.* [SXZ\*12] adopt an interactive approach to labelled segmentations of RGBD images, followed by 3D model retrieval for scene modelling. Chen *et al.* [CLW\*14] learn contextual relationships from a 3D scene database to further constrain the reconstruction for semantic compatibility between both objects and parts.

In contrast to considering entire 3D shapes, Gal *et al.* [GSH\*07] utilize a dictionary of local shape priors (defined as patches) for shape reconstruction. The method is mainly designed for enhancing shape features, where each region of an input point cloud is matched to a shape patch in the database. The matched shape patch is then used to enhance and rectify the local region. Recently, Mattausch *et al.* [MPM\*14] introduce a patch-based reconstruction system for indoor scenes. Their method considers recognizing and fitting planar patches from point cloud data.

Shen *et al.* [SFCH12] extends this idea for single object reconstruction by assembling object parts. Their method utilizes a collection consistently segmented 3D shapes. Given a scan of an object, the method recursively searches for parts in the collection to assemble the original object. The retrieval phase considers both the geometric similarity between the input and retrieved parts as well as part compatibility which is learned from the input shapes. Sung *et al.* [SKAG15] describe a framework for reliably estimating the part structure of partial scans by treating each part as a local coordinate system. Their method also utilizes symmetric properties of the target object and shape collection, providing more accurate reconstructions on their shape completion benchmark.

**Data-driven SLAM.** Non-parametric methods have also found applications in reconstructing temporal geometric data (e.g. the output of the Kinect scanner). The simultaneous localization and mapping (or SLAM) method is a notable technique which jointly estimates the trajectory of the scanning device alongside the geometry of the environment. In this case, shape collections serve as priors for the objects in the environment, which can be used to train object detectors. For example, the SLAM++ system proposed

by Salas-Moreno *et al.* [SMNS\*13] trained domain-specific object detectors from shape collections. The learned detectors are integrated inside the SLAM framework to recognize and track those objects. Similarly, Kim *et al.* [KMYG12] use learned object models to reconstruct dense 3D models from a single scan of an indoor scene. More recently, Song *et al.* [SX14] introduced a 3D sliding window object detector with improved performance and capability extending to a broader range of objects. Li *et al.* [LDGN15] propose a data-assisted online reconstruction technique which allows object retrieval from a 3D shape database while simultaneously scanning an environment in real-time.

**Shape-driven reconstruction from images.** Recently, there is a growing interest in reconstructing 3D objects directly from images (e.g. [XZZ\*11, KSES14, AME\*14, SHM\*14]). This problem introduces fundamental challenges in both querying similar objects and deforming objects/parts to fit the input object. In terms of searching similar objects, successful methods typically render objects in the database from a dense set of viewpoints and pick objects where one view is similar to the input image object. Since depth information is missing from the image object, it is important to properly regularize 3D object transformations; otherwise, a 3D object may be deformed arbitrarily even though its projection on the image domain matches the image object. Most existing techniques consider rigid transformations or user-specified deformations [XZZ\*11]. In a recent work, Su *et al.* [SHM\*14] propose to learn meaningful deformations of each shape from its optimal deformation to similar shapes. Huang *et al.* [HWK15] jointly analyse a large collection of images in object categories and a smaller collection of 3D models to achieve simultaneous analysis and reconstruction of 2D images.

## 7. Shape Modelling and Synthesis

So far, the creation of detailed three-dimensional content remains a tedious task which can mainly be performed by skilled artists. Three-dimensional content creation has been a major bottleneck hindering the development of ubiquitous 3D graphics. Thus, providing easy-to-use tools for casual and novice users to design and create 3D models has been a key challenge in computer graphics. To address this challenge, current literature has been focused on two main directions, i.e. intelligent interfaces for interactive shape modelling and smart models for automated model synthesis. The former strives to endow modelling interfaces with a higher level understanding of the structure and semantics of 3D shapes, allowing the interface to reason around the incomplete shapes being modelled. The latter


direction focuses on developing data-driven models to synthesize new shapes automatically. The core problem is to learn generative shape models from a set of exemplars (e.g. probability distributions, fitness functions, functional constraints, etc.) so that the synthesized shapes are plausible and novel. It can be seen that both of the two paradigms depend on data-driven modelling of shape structures and semantics. With the availability of large 3D shape collections, the data-driven approach becomes a promising answer to the content creation bottleneck.

### 7.1. Interactive shape modelling and editing

Interactive 3D modelling software (3DS Max, Maya, etc.) provide artists with a large set of tools for creating and editing detailed 3D models. Unfortunately, this same software is often onerous to harness for non-professional users. For casual users, more intuitive modelling interfaces with a certain machine intelligence are to be preferred. Below, we discuss such methods for assembly-based modelling and guided shape editing.

**Assembly-based modelling.** Early works on 3D modelling based on shape sets are primarily driven by the purpose of *content reuse* in part-assembly based modelling approaches. The seminal work of modelling by example [FKS\*04] presents a pioneering system of shape modelling by searching a shape database for parts to reuse in the construction of new shapes. Kraevoy *et al.* [KJS07] describe a system for shape creation via interchanging parts between a small set of compatible shapes. Guo *et al.* [GLXJ14] propose assembly-based creature modelling guided by a shape grammar.

Beyond content reuse through database queries or hand-crafted rules, Chaudhuri and Koltun [CK10] propose a data-driven technique for suggesting the modeller with shape parts that can potentially augment the current shape being built. Such part suggestions are generated through retrieving a shape from a database based on partial shape matching. Although this is a purely geometric method without accounting for the semantics of shape parts, it represents the first attempt at utilizing a shape database to *augment the modelling interface*. Later, Chaudhuri *et al.* [CKGK11] show that the incorporation of semantic relationships increases the relevance of presented parts. Given a repository of 3D shapes, the method learns a probabilistic graphical model encoding semantic and geometric relationships among shape parts. During modelling, inference in the learned Bayesian network is performed to produce a relevance ranking of the parts (see Figure 13).



**Figure 13:** Given a library of models, a Bayesian network encoding semantic and geometric relationships among shape parts is learned [CKGK11] (top). The modelling process (bottom) performs probabilistic inference in the learned Bayesian network to generate ranked lists of category labels and components within each category, customized for the currently assembled model.



A common limitation of the above techniques is that they do not provide a way to directly express a high-level design goal (e.g. ‘create a cute toy’). Chaudhuri *et al.* [CKG\*13] proposed a method that learns semantic attributes for shape parts that reflect the high-level intent people may have for creating content in a domain (e.g. adjectives such as ‘dangerous’, ‘scary’ or ‘strong’) and ranks them according to the strength of each learned attribute (Figure 5). During an interactive session, the user explores and modifies the strengths of semantic attributes to generate new part assemblies.

Three-dimensional shape collections can supply other useful information as well, such as contextual and spatial relationships between shape parts, to enhance a variety of modelling interfaces. Xie *et al.* [XXM\*13] propose a data-driven sketch-based 3D modelling system. In the offline learning stage, a shape database is pre-analysed to extract the contextual information among parts. During the online stage, the user designs a 3D model by progressively sketching its parts and retrieving and assembling shape parts from the database. Both the retrieval and assembly are assisted by pre-computed contextual information so that more relevant parts can be returned and selected parts can be automatically placed. Inspired by the Shadow-Draw system [LZC11], Fan *et al.* [FWX\*13] propose 3D modelling by drawing with data-driven shadow guidance. The user’s strokes are used to query a 3D shape database for generating the shadow image, which in turn can guide the user’s drawing. Along the drawing, 3D candidate parts are retrieved for assembly-based modelling. Starting from a collection of expertly created, fabricable 3D models, Schulz *et al.* [SSL\*14] extract parameterized design templates encoding all information necessary for fabrication. The templates can then be used to generate new fabricable models in an interactive design system.

**Shape editing.** The general idea of data-driven shape editing is to learn a model from a collection of closely related shapes that characterizes the plausible variations or deformations of the shapes in this collection. In this way, the learned model can be used to constrain a user’s edit to maintain plausibility. For organic shapes, such as human faces [BV99, CWZ\*14] or bodies [ACP03], parametric models can be learned from a shape set characterizing its shape space. Such parametric models can be used to edit the shapes through exploring the shape space with the set of parameters.

An alternative widely adopted approach is the analyse-and-edit paradigm. This technique first extracts the structure of the input shape, and then uses this structure to constrain the editing phase to be more tenable [GSMCO09]. Instead of learning structure from a single shape, Fish *et al.* [FAvK\*14] learn it from a set of shapes that belong to the same family, resulting in a set of probability distributions characterizing the part arrangements. These distributions can be used to guide structure-preserving editing, where models can be edited while maintaining their familial traits. Yumer and Kara [YK14] extract co-constrained handles from a set of shapes for shape deformation. The handles are generated based on co-abstraction [YK12] of the set of shapes and the deformation co-constraints are learned statistically from the set. The deformation handles can also be controlled with continuous semantic attributes [YCHK15].

Based on learned structures from a database of 3D models, Xu *et al.* [XZZ\*11] propose photo-inspired 3D object modelling.


Guided by the object in a photograph, the method creates a 3D model as a geometric variation of a candidate model retrieved from the database. Due to the pre-analysed structural information, the method addresses the ill-posed problem of 3D modelling from a single 2D image via structure-preserving 3D warping. The final result is structurally plausible and is readily usable for subsequent editing. Moreover, the resulting 3D model, although built from a single view, is structurally coherent from all views.

## 7.2. Automated synthesis of shapes

Many applications such as 3D games and films require large collections of 3D shapes for populating their environments. Modelling each shape individually can be tedious even with the best interactive tools. The goal of data-driven shape synthesis algorithms is to generate several shapes automatically with no or very little user supervision: users may only provide some preferences or high-level specifications to control the shape synthesis process. Existing methods achieve this task by using probabilistic generative models of shapes, evolutionary methods or learned probabilistic grammars.

**Statistical models of shapes.** The basic idea of these methods is to define a parametric shape space and then fit a probability distribution to the data points that represent the input exemplar shapes. Since the input shapes are assumed to be plausible and desired representatives of the shape space, high-probability areas of the shape space which tend to become associated with new, plausible shape variants. This idea was first explored in the context of parametric models [BV99, ACP03], discussed in Section 6. By associating each principal component of the shape space defined by these methods with a Gaussian distribution, this distribution can be sampled to generate new human faces or bodies (Figure 11). Since the probability distribution of plausible shapes tends to be highly non-uniform in several shape classes, Talton *et al.* [TGY\*09] use kernel density estimation with Gaussian kernels to represent plausible shape variability. The method is able to generate new shapes for tree and human body parametric spaces.

Shapes have structure, i.e. shapes vary in terms of their type and style, different shape styles have different number and type of parts, parts have various subparts that can be made of patches, and so on. Thus, to generate shapes in complex domains, it is important to define shape spaces over structural and geometric parameters, and capture hierarchical relationships between these parameters at different levels. Kalogerakis *et al.* [KCKK12] (Figure 14) proposed a



**Figure 14:** Given a hundred training aeroplanes (in green), the probabilistic model from [KCKK12] synthesizes several hundreds of new aeroplanes (in blue).

probabilistic model that represents variation and relationships of geometric descriptors and adjacency features for different part styles, as well as variation and relationships of part styles and repetitions for different shape styles. The method learns the model from a set of consistently segmented shapes. Part and shape styles are discovered based on latent variables that capture the underlying modes of shape variability. The method uses a search procedure to assemble new shapes from parts of the input shapes according to the learned probability distribution. Users can also set preferences for generating shapes from a particular shape style, with given part styles or specific parts. Instead of relying on pre-segmented shapes and high-level part descriptors to encode shape variability, Huang *et al.* [HKM15] propose a probabilistic model that jointly estimates shape segmentation, surface correspondences and surface descriptors from an input shape data set. A deep learning procedure was used to capture hierarchical relationships of corresponding surface point positions and parts as well as their existence in the input shapes. Their probabilistic model can be sampled to directly generate point-sampled surface geometry and shape structure.

**Set evolution.** Xu *et al.* [XZCOC12] developed a method for generating shapes inspired by the theory of evolution in biology. The basic idea of set evolution is to define cross-over and mutation operators on shapes to perform part warping and part replacement. Starting from an initial generation of shapes with part correspondences and built-in structural information such as interpart symmetries, these operators are applied to create a new generation of shapes. A selected subset from the generation is presented via a gallery to the user who provides feedback to the system by rating them. The ratings are used to define the fitness function for the evolution. Through the evolution, the set is personalized and populated with shapes that better fit to the user. At the same time, the system explicitly maintains the diversity of the population so as to prevent it from converging into an ‘elite’ set.

**Learned shape grammars.** Talton *et al.* [TYK\*12] leverage techniques from natural language processing to learn probabilistic generative grammars of shapes. The method takes as input a set of exemplar shapes represented with a scene graph specifying parent/child relationships and relative transformations between labelled shape components. They use Bayesian inference to learn a probabilistic formal grammar that can be used to synthesize novel shapes.

## 8. Scene Analysis and Synthesis

Analysing and modelling indoor and outdoor environments has important applications in various domains. For example, in robotics it is desirable for an autonomous agent to understand the semantics of 3D environments to be able to interact with them. In urban planning and architecture, professionals build digital models of cities and buildings to validate and improve their designs. In computer graphics, artists create novel 3D scenes for movies and video games.

The fast growing number of 3D scenes in digital repositories provide new opportunities for data-driven scene analysis, editing and synthesis. Emerging collections of 3D scenes pose novel research challenges that cannot be easily addressed with existing tools. In particular, representations created for analysing collections of single models mostly focus on arrangement and relations between shape

parts [MWZ\*14], which usually exhibit less variations than objects in scenes. Capturing scene structure poses a greater challenge due to looser spatial relations and a more diverse mixture of functional substructures.

Inferring scene semantics is a long-standing problem in image understanding, with many methods developed for object recognition [QT09], classification [SW10], layout and structure reasoning [CCPS13, FGH13] with a *single image*. Previous work demonstrates that one can leverage collections of 3D models to facilitate scene understanding in images [SLH12]. In addition, the depth information in RGBD scans can be used to establish the link between 2D and 3D for model-driven scene understanding [SKHF12]. The semantic annotations of images are not immediately useful for modelling and synthesizing 3D scenes, for which the geometric and structural priors have to be learned from 3D data.


In this section, we cover the data-driven techniques that leverage collections of 3D scenes for modelling, editing and synthesizing novel scenes.

**Context-based retrieval.** To address the large variation in the geometry and arrangement of objects in scenes, Fisher *et al.* [FH10, FSH11] propose to take advantage of local context. One of the key insights of their work is that collections of 3D scenes provide rich information about context in which objects appear. They show that capturing these contextual priors can help in scene retrieval and editing.


Their system takes an annotated collection of 3D scenes as input, where each object in a scene is classified. They represent each scene as a graph, where nodes represent objects and edges represent relations between objects, such as support and surface contact. In order to compare scenes, they define kernel functions for pairs of nodes measuring similarity in object geometry, and for pairs of edges, measuring similarity in relations of two pairs of objects. They further define a graph kernel to compare pairs of scenes. In particular, they compare all walks of fixed length originating at all pairs of objects in both scene graphs, which loosely captures similarities of all contexts in which objects appear [FSH11]. They show that this similarity metric can be used to retrieve scenes. By comparing only paths originated at a particular object, they can retrieve objects for interactive scene editing.

**Focal points.** Measuring the similarity of complex hybrid scenes such as studios composed of a bedroom, living room and dining room poses a challenge to graph kernel techniques since they only measure global scene similarity. Thus, Xu *et al.* [XMZ\*14] advocate analysing salient sub-scenes, which they call focal points, to compare hybrid scenes, i.e. scenes containing multiple salient sub-scenes. Figure 15 shows an example of comparing complex scenes, where the middle scene is a hybrid one encompassing two semantically salient sub-scenes, i.e. bed-nightstands and TV-table-sofa. The middle scene is closer to the left one when the bed and nightstands are focused on, and otherwise when the TV-table-sofa combo is the focal point. Therefore, scene comparison may yield different similarity distances depending on the focal points.

Formally, a focal point is defined as a representative substructure of a scene which can characterize a semantic scene category. That



**Figure 15:** Scene comparisons may yield different similarity distances (left) depending on the focal points [XMZ\*14].



**Figure 16:** The algorithm processes raw scene graphs with possible over-segmentation (a) into consistent hierarchies capturing semantic and functional groups (b,c) [LCK\*14].

means the substructure should reoccur frequently only within that category. Therefore, focal point detection is naturally coupled with the identification of scene categories via scene clustering. This poses coupled problems of detecting focal points based on scene groups and grouping scenes based on focal points. These two problems are solved via interleaved optimization which alternates between focal point detection and focal-based scene clustering. The former is achieved by mining frequent substructures and the latter uses subspace clustering, where scene distances are defined in a focal-centric manner. Inspired by the work of Fisher *et al.* [FSH11], scene distances are computed using focal-centric graph kernels which are estimated from walks originating from representative focal points.


The detected focal points can be used to organize the scene collection and to support efficient exploration of the collection (see Section 9). Focal-based scene similarity can be used for novel applications such as multi-query scene retrieval, where one may issue queries consisting of multiple semantically related scenes and wish to retrieve more scenes ‘of the same kind’.

**Synthesis.** Given an annotated scene collection, one can also synthesize new scenes that have a similar distribution of objects. The scene synthesis technique of Fisher *et al.* [FRS\*12] learns two probabilistic models from the training data set: (1) object occurrence, indicating which objects should be placed in the scene, and (2) layout optimization, indicating where to place the objects. Next, it takes an example scene, and then synthesizes similar scenes using the learned priors. It replaces or adds new objects using context-based retrieval techniques, and then optimizes for object placement based on learned object-to-object spatial relations. Synthesizing example scenes might be a challenging task, thus Xu *et al.* [XCF\*13] propose modelling 3D indoor scenes from 2D sketches, by lever-

aging a database of 3D scenes. Their system jointly optimizes for sketch-guided co-retrieval and co-placement of all objects.

**Hierarchical scene annotation.** All aforementioned applications take an annotated collection of 3D scenes as an input. Unfortunately, most scenes in public repositories are not annotated and thus require additional manual labelling [FRS\*12]. Liu *et al.* [LCK\*14] address the challenge of annotating novel scenes. The key observation of their work is that understanding hierarchical structure of a scene enables efficient encoding of functional scene substructures, which significantly simplifies detecting objects and representing their relationships. Thus, they propose a supervised learning approach to estimate a hierarchical structure for novel scenes. Given a collection of scene graphs with consistent hierarchies and labels, they train a probabilistic hierarchical grammar encoding the distributions of shapes, cardinalities and spatial relationships between objects. Such a grammar can then be used to parse new scenes: find segmentations, object labels and hierarchical organization of objects consistent with the annotated collection (see Figure 16).

**Challenges and opportunities.** The topic of 3D scene analysis is quite new and there are many open problems and research opportunities. The first problem is to efficiently characterize spatial relationships between objects and object groups. Most existing methods work with bounding box representation which are efficient to process, but not sufficiently informative to characterize object-to-object relationships. For example, one cannot reliably determine the object enclosure relationship based on a bounding box. Recently, Zhao *et al.* [ZWK14] propose to use biologically inspired bisector surface to characterize the geometric interaction between adjacent objects for 3D scene indexing (Figure 17). The bisector surface can be



**Figure 17:** The interaction bisector surface (in blue) of several two-object scenes [ZWK14].

extended into a geometric descriptor for contextual modelling of the functionality of a 3D object in a given scene [HZvK\*15]. Second, most existing techniques heavily rely on expert user supervision for scene understanding. Unfortunately, online repositories rarely have models with reliable object tags. Therefore, there is a need for methods that can leverage scenes containing only partial and/or noisy annotations. Finally, the popularity of commodity RGBD cameras has significantly simplified the acquisition of indoor scenes. This emerging scanning technique opens space for new applications such as online scene analysis [XTZ14, XHS\*15].

## 9. Exploration and Organization

The rapidly growing quantity and variety of digital 3D models in large online collections have caused an emerging need to develop algorithms and techniques that effectively organize these large collections and allow users to interactively explore them. For example, an architect might furnish a digital building by searching through databases organized according to furniture types, regions of interest and design styles. Likewise, an industrial designer can explore shape variations among existing products when creating a new object. Most existing repositories only support text-based search, relying on user-entered tags and titles. This approach suffers from inaccurate and ambiguous tags, often entered in different languages. While it is possible to try using shape analysis to infer consistent tags as discussed in Section 3, it is difficult to convey stylistic and geometric variations using only text. An alternative approach can be to perform shape, sketch or image-based queries. However, to formulate such search queries the user needs to have a clear mental model of the shape that should be retrieved. Thus, some researchers focus on providing tools for *exploring* shape collections. Unlike search, exploration techniques do not assume *a priori* knowledge of the repository content, and help the user to understand geometric, topological and semantic variations within the collection.

**Problem statement and method categorization.** Data exploration and organization is a classical problem in data analysis and visualization [PEP\*11]. Given a data collection, the research focuses on *grouping and relating data points, learning the data variations in the collection and organizing the collection into a structured form*, to facilitate retrieval, browsing, summarization and visualization of the data, based on efficient *interfaces or metaphors*.

The first step to organizing model collections is to devise appropriate metrics to relate different data points. Various similarity metrics have been proposed in the past to relate entire shapes as well as local regions on shapes. In particular, previous sections of this document cover algorithms for computing global shape similarities

**Table 4:** A summary of several recent works over four aspects. *Metaphor:* templates, surface painted ROIs, probability distribution plots or query shapes. *Shape Comparison:* shape similarity, part or point correspondence. *Variability:* geometry, topology or both. *Organization Form:* cluster or hierarchy.

| Method    | Meta. | Comp. | Var.  | Org.      |
|-----------|-------|-------|-------|-----------|
| [OLGM11]  | temp. | simi. | geom. | n/a       |
| [KLM*13]  | temp. | part  | both  | cluster   |
| [AKZM14]  | plot  | part  | both  | cluster   |
| [KLM*12]  | ROI   | point | both  | n/a       |
| [ROA*13]  | ROI   | point | geom. | n/a       |
| [HWG14]   | ROI   | point | both  | cluster   |
| [XMZ*14]  | ROI   | simi. | topo. | cluster   |
| [FAvK*14] | plot  | part  | geom. | cluster   |
| [HSS*13]  | query | simi. | both  | hierarchy |

(Section 3), partwise correspondences (Section 4) and pointwise correspondences (Section 5). In this section, we will focus on techniques that take advantage of these correlations to provide different interfaces for exploring and understanding geometric variability in collections of 3D shapes. We categorize the existing exploration approaches based on four aspects:


- **Metaphor:** a user interface for exploring shape variations. We will discuss five basic exploration interfaces, ones that use proxy shapes (templates), regions of interest, probability plots, query shapes or continuous attributes.
- **Shape comparison:** techniques used to relate different shapes. We will discuss techniques that use global shape similarities, as well as part or point correspondences.
- **Variability:** shape variations captured by the system. Most methods we will discuss rely on geometric variability of shapes or parts. Some techniques also take advantage of topological variability; that is, variance in number of parts or how they are connected (or variance in numbers of objects and their arrangements in scenes).
- **Organizational form:** a method to group shapes. We will discuss methods that group similar shapes to facilitate exploring intra-group similarities and intergroup variations, typically including clustering and hierarchical clustering.

Table 4 summarizes several representative works in terms of these aspects. In the remaining part of this section, we list several recent techniques which are grouped based on the exploration metaphor.

**Template-based exploration.** Componentwise variability in position and scale of parts reveals useful information about a model collection. Several techniques use box-like templates to show variations among models of the same class. Ovsjanikov *et al.* [OLGM11] describe a technique for learning these partwise variations without solving the challenging problem of consistent segmentation. First, they use the segmentation of a single shape to construct the initial template. This is the only step that needs to be verified and potentially fixed by the user. The next goal is to automatically infer deformations of the template that would capture the most important geometric variations of other models in the

**Table 5:** Comparison of various works on data-driven shape analysis and processing. For each work, we summarize over the criterion set defined for data-driven methods: training data (encompassing data representation, preprocessing and scale), feature (including feature type and whether feature selection is involved), learning model or approach, learning type (supervised, semi-supervised and unsupervised), learning outcome (e.g. a classifier or a distance metric), as well as its typical application scenario. See the text for detailed explanation of the criteria. Some works employ another work as a pre-processing stage (e.g. [CKG\*13] requires the labelled segmentation produced by [KHS10]). There are four types of features including local geometric features (Local), global shape descriptors (Global), both local and global shape features (L.&G.), structural features (Struct.) as well as 2D visual features (Visual).

| Work      | Training data |                |           | Feature |      | Learning model/approach         | Learning type   | Learning outcome      | Application    |
|-----------|---------------|----------------|-----------|---------|------|---------------------------------|-----------------|-----------------------|----------------|
|           | Rep.          | Pre-proc.      | Scale     | Type    | Sel. |                                 |                 |                       |                |
| [FKMS05]  | Point         | No             | Thousands | Local   | No   | SVM classifier                  | Supervised      | Object classifier     | Classification |
| [BBOG11]  | Mesh          | No             | Thousands | Local   | No   | Similarity Sensitive Hashing    | Supervised      | Distance metric       | Classification |
| [HSG13]   | Mesh          | Pre-align.     | Thousands | Local   | No   | Max-marginal distance learning  | Semi-supervised | Distance metric       | Classification |
| [KHS10]   | Mesh          | No             | Tens      | Local   | Yes  | Jointboost classifier           | Supervised      | Face classifier       | Segmentation   |
| [vKTS*11] | Mesh          | Yes            | Tens      | Local   | Yes  | Gentleboost classifier          | Supervised      | Face classifier       | Segmentation   |
| [BLVD11]  | Mesh          | No             | Tens      | L.&G.   | Yes  | Adaboost classifier             | Supervised      | Boundary classifier   | Segmentation   |
| [XXLX14]  | Mesh          | No             | Hundreds  | Local   | Yes  | Feedforward neural networks     | Supervised      | Face/patch classifier | Segmentation   |
| [XSX*14]  | Mesh          | Pre-seg.       | Tens      | Local   | No   | Sparse model selection          | Supervised      | Segment similarity    | Segmentation   |
| [LCHB12]  | Mesh          | No             | Tens      | Local   | Yes  | Entropy regularization          | Semi-supervised | Face classifier       | Segmentation   |
| [WAVK*12] | Mesh          | Pre-seg.       | Hundreds  | Local   | No   | Active learning                 | Semi-supervised | Segment classifier    | Segmentation   |
| [WGW*13]  | Image         | Labelled parts | Hundreds  | Local   | No   | 2D shape matching               | Supervised      | 2D shape similarity   | Segmentation   |
| [HFL12]   | Mesh          | Over-seg.      | Tens      | Local   | Yes  | Subspace clustering             | Unsupervised    | Patch similarity      | Seg./Corr.     |
| [SvKK*11] | Mesh          | Pre-seg.       | Tens      | Local   | No   | Spectral clustering             | Unsupervised    | Seg. simi./classifier | Seg./Corr.     |
| [XLZ*10]  | Mesh          | Part           | Tens      | Struct. | No   | Spectral clustering             | Unsupervised    | Part proportion simi. | Seg./Corr.     |
| [vKXZ*13] | Mesh          | Part           | Tens      | Struct. | No   | Multi-instance clustering       | Unsupervised    | Seg. hier. simi.      | Seg./Corr.     |
| [GF09]    | Mesh          | No             | Tens      | Global  | No   | Global shape alignment          | Unsupervised    | Face similarity       | Seg./Corr.     |
| [HKG11]   | Mesh          | Pre-seg.       | Tens      | Local   | No   | Joint part matching             | Unsupervised    | Segment similarity    | Seg./Corr.     |
| [HWG14]   | Mesh          | Init. corr.    | Tens      | Global  | No   | Consistent func. map networks   | Unsupervised    | Segment similarity    | Seg./Corr.     |
| [KLM*13]  | Mesh          | Template       | Thousands | Local   | No   | Shape alignment                 | Semi-supervised | Templates             | Seg./Corr.     |
| [MPM*14]  | Mesh          | Over-seg.      | Hundreds  | Local   | No   | Density-based clustering        | Unsupervised    | Patch similarity      | Recognition    |
| [NBCW*11] | Mesh          | Init. corr.    | Tens      | L.&G.   | No   | Inconsistent map detection      | Unsupervised    | Point similarity      | Corr./Expl.    |
| [HZG*12]  | Mesh          | Init. corr.    | Tens      | L.&G.   | No   | MRF joint matching              | Unsupervised    | Point similarity      | Corr./Expl.    |
| [KLM*12]  | Mesh          | Pre-align.     | Tens      | Global  | No   | Spectral matrix recovery        | Unsupervised    | Point similarity      | Corr./Expl.    |
| [HG13]    | Mesh          | Init. corr.    | Tens      | Global  | No   | Low-rank matrix recovery        | Unsupervised    | Point similarity      | Corr./Expl.    |
| [OLGM11]  | Mesh          | Part           | Hundreds  | Global  | No   | Manifold learning               | Unsupervised    | Parametric model      | Exploration    |
| [ROA*13]  | Mesh          | Map            | Tens      | None    | N/A  | Functional map analysis         | Unsupervised    | Difference operator   | Exploration    |
| [FAVK*14] | Mesh          | Labelled parts | Hundreds  | Struct. | No   | Kernel Density Estimation       | Supervised      | Prob. distributions   | Expl./Synth.   |
| [AKZM14]  | Mesh          | [KLM*13]       | Thousands | Struct. | No   | Manifold learning               | Unsupervised    | Parametric models     | Expl./Synth.   |
| [HSS*13]  | Mesh          | No             | Hundreds  | Global  | No   | Quartet analysis and clustering | Unsupervised    | Distance measure      | Organization   |
| [BV99]    | Mesh          | Pre-align.     | Hundreds  | Local   | No   | Principal Component Analysis    | Unsupervised    | Parametric model      | Recon./Expl.   |
| [ACP03]   | Point         | Pre-align.     | Hundreds  | Local   | No   | Principal Component Analysis    | Unsupervised    | Parametric model      | Recon./Expl.   |
| [HSS*09]  | Point         | Pre-align.     | Hundreds  | Local   | No   | PCA & linear regression         | Unsupervised    | Parametric model      | Recon./Expl.   |
| [PMG*05]  | Mesh          | Pre-align.     | Hundreds  | Global  | No   | Global shape alignment          | Unsupervised    | Shape similarity      | Reconstruction |
| [NXS12]   | Point         | Labelled parts | Hundreds  | Struct. | No   | Random Forest Classifier        | Supervised      | Object classifier     | Reconstruction |
| [SFCH12]  | Mesh          | Labelled parts | Tens      | Global  | No   | Part matching                   | Unsupervised    | Part detector         | Reconstruction |
| [KMYG12]  | Point         | Labelled parts | Tens      | Local   | No   | Joint part fitting and matching | Unsupervised    | Object detector       | Reconstruction |
| [SMNS*13] | Mesh          | No             | Tens      | L.&G.   | No   | Shape matching                  | Unsupervised    | Object detector       | Reconstruction |
| [XZZ*11]  | Mesh          | Labelled parts | Tens      | Struct. | No   | Structural shape matching       | Unsupervised    | Part detector         | Modelling      |
| [AME*14]  | Mesh          | Projected      | Thousands | Visual  | No   | Linear Discriminant Analysis    | Supervised      | Object detector       | Recognition    |
| [SHM*14]  | Mesh          | Projected      | Tens      | Visual  | No   | Shape matching                  | Unsupervised    | 2D-3D correlation     | Reconstruction |
| [CK10]    | Mesh          | No             | Thousands | Global  | No   | Shape matching                  | Unsupervised    | Part detector         | Modelling      |
| [CKGK11]  | Mesh          | [KHS10]        | Hundreds  | Local   | No   | Bayesian Network                | Unsupervised    | Part reasoning model  | Modelling      |
| [XXM*13]  | Mesh          | Labelled parts | Tens      | Struct. | No   | Contextual part matching        | Unsupervised    | Part detector         | Modelling      |
| [KCKK12]  | Mesh          | [KHS10]        | Hundreds  | L.&G.   | No   | Bayesian Network                | Unsupervised    | Shape reasoning model | Synthesis      |
| [XZCOC12] | Mesh          | Part           | Tens      | Struct. | No   | Part matching                   | Unsupervised    | Part similarity       | Synthesis      |
| [TYK*12]  | Mesh          | Labelled parts | Tens      | Struct. | No   | Structured concept learning     | Unsupervised    | Probabilistic grammar | Synthesis      |
| [YK12]    | Mesh          | No             | Tens      | Global  | No   | Shape matching                  | Unsupervised    | Shape abs. similarity | Modelling      |
| [YK14]    | Mesh          | Pre-seg.       | Tens      | Local   | No   | Segment matching                | Unsupervised    | Segment abs. simi.    | Modelling      |
| [CKG*13]  | Mesh          | [KHS10]        | Hundreds  | L.&G.   | No   | SVM ranking                     | Supervised      | Ranking metric        | Model/Expl.    |
| [FSH11]   | Scene         | Labelled obj.  | Tens      | Struct. | No   | Relevance feedback              | Supervised      | Contextual obj. simi. | Classification |
| [FRS*12]  | Scene         | Labelled obj.  | Hundreds  | Struct. | No   | Bayesian Network                | Supervised      | Mixture models        | Synthesis      |
| [XCF*13]  | Scene         | Labelled obj.  | Hundreds  | Struct. | No   | Frequent subgraph mining        | Unsupervised    | Frequent obj. groups  | Modelling      |
| [XMZ*14]  | Scene         | Labelled obj.  | Hundreds  | Struct. | No   | Weighted subgraph mining        | Unsupervised    | Distinct obj. groups  | Org./Expl.     |
| [LCK*14]  | Scene         | Labelled hier. | Tens      | Struct. | No   | Probabilistic learning          | Supervised      | Probabilistic grammar | Seg./Corr.     |



**Figure 18:** Shape exploration based on fuzzy correspondence. The user paints a region of interest (ROI) on a query shape (left column), and the method sorts models based on their similarity within the region (right).

collection. They hypothesize that all shapes can be projected on a low-dimensional manifold based on their global shape descriptors. Finally, they reveal the manifold structure by deforming a template to fit to the sample points. Directions for interesting variations are depicted by arrows on the template and the shapes that correspond to the current template configuration are presented to the user.

The descriptor-based approach described above assumes that all intraclass shapes share the same parts and that there exists a low-dimensional manifold that can be captured by deforming a single template. These assumptions do not hold for large and diverse collections of 3D models. To tackle this challenge, Kim *et al.* [KLM\*13] proposed an algorithm for learning several part-based templates capturing multi-modal variability in collections of shapes. They start with an initial template that includes a superset of all parts that might occur in a data set, and jointly learn part segmentations, point-to-point surface correspondence as well as a compact deformation model. The output is a set of templates that groups the input models into clusters, capturing their styles and variations.


**ROI-based exploration.** Not all interesting variations occur at the scale of parts: they can occur at subpart scale, or span multiple subregions from multiple parts. In these cases, the user may prefer to select an arbitrary region on a 3D model and look for more models sharing similar regions of interest. Such detailed and flexible queries require a finer understanding of correspondences between different shapes. Kim *et al.* [KLM\*12] propose fuzzy point correspondences to encode the inherent ambiguity in relating diverse shapes. Fuzzy point correspondences are represented by real values specified for all pairs of points, indicating how well the points correspond. They leverage transitivity in correspondence relationships to compute this representation from a sparse set of pairwise point correspondences. The interface proposed by Kim *et al.* allows users to paint regions of interest directly on a surface and then retrieve similar regions among other shapes, or even show geometric variations found in the selected region (see Figure 18).

One limitation of correspondence-based techniques is that they typically do not consider the entire collection when estimating shape differences. Rustamov *et al.* [ROA\*13] focus on a fundamental intrinsic representation for shape differences. Starting with a functional map between two shapes, that is, a map that describes a


change of functional basis, they derive a shape difference operator revealing detailed information about the location, type and magnitude of distortions induced by a map. This makes shape difference a quantifiable object that can be co-analysed within a context of the entire collection. They show that this deeper understanding of shape differences can help in exploration. For example, one can embed shapes in a low-dimensional space based on shape differences, or use shape difference to interpolate variations by showing ‘intermediate’ shapes between two regions of interest. To extend these techniques to man-made objects, Huang *et al.* [HWG14] construct a consistent functional basis for shape collections exhibiting large geometric and topological variability. They show that the resulting consistent maps are capable of capturing discrete topological variability, such as variance in the number of bars of the back of a chair.

**ROI-based scene exploration.** Recent works on organizing and exploring 3D visual data mostly focus on object collections. Exploring 3D scenes poses additional challenges since scenes typically exhibit more structural variations. Unlike man-made objects that usually contain a handful of object parts, scenes can contain anywhere from ten to hundreds of objects. Not only this, but the objects themselves do not typically have a prescribed rigid arrangement with respect to each other. Thus, global scene similarity metrics, such as the graph kernel based one used in [FRS\*12] are limited to organizing data sets based on very high level features, such as scene type. Xu *et al.* [XMZ\*14] advocate that 3D scenes should be compared from the perspective of a particular focal point which is a representative substructure of a specific scene category. Focal points are detected through contextual analysis of a collection of scenes, resulting in a clustering of the scene collection where each cluster is characterized by its representative focal points (see Section 8). Consequently, the focal points extracted from a scene collection can be used to organize collection into an interlinked and well-connected cluster formation, which facilitates scene exploration. Figure 19 shows an illustration of such cluster-based organization and an exploratory path transiting between two scene clusters/categories.

**Plot-based exploration.** All aforementioned exploration techniques typically do not visualize the probabilistic nature of shape variations. Fish *et al.* [FAvK\*14] study the configurations of shape parts from a probabilistic perspective, trying to indicate which shape variations are more likely to occur. To learn the distributions of part arrangements, all shapes in the family are pre-segmented consistently. The resulting set of probability density functions (PDFs) characterizes the variability of relations and arrangements across different parts. A peak in a PDF curve represents that particular a configuration of the related parts frequently appeared among several shapes in the family. The multiple PDFs can be used as interfaces to interactively explore the shape family from various perspectives. Averkiou *et al.* [AKZM14] use part structure inferred by this method to produce a low-dimensional part-aware embedding of all models. The user can explore interesting variations in part arrangements simply by moving the mouse over the 2D embedding. In addition, their technique allows the synthesis of novel shapes by clicking on empty spaces in the embedded space. Upon clicking, the system would deform parts from neighbouring shapes to synthesize a novel part arrangement.



**Figure 19:** Focal-based scene clustering produces overlapping clusters, which is due to hybrid scenes possessing multiple focal points. An exploratory path, from (a) to (e), through the overlap, smoothly transit between the two scene clusters, representing bedroom and offices, respectively.



**Figure 20:** Given a set of heterogeneous shapes, a reliable qualitative similarity is derived from quartets composed of two pairs of objects (left). Aggregating such qualitative information from many quartets computed across the whole set leads to a categorization tree as a hierarchical organization of the input shape collection (right).

**Query-based exploration** For a heterogeneous shape collection encompassing diverse object classes, it is typically not possible to characterize part-structure and correspondences between all pairs of shapes. Even global shape similarity is not a very reliable feature in this setting, which makes organizing and exploring heterogeneous collections especially difficult. To address this challenge, Huang *et al.* [HSS\*13] introduce qualitative analysis techniques from the field of bioinformatics. Instead of relying on quantitative distances, which may be ill-applied between dissimilar shapes, the method considers a more reliable *qualitative similarity* derived from *quartets* composed of two pairs of objects. The shapes that are paired in the quartet are close to each other and far from the shapes in the other pair, where distances are estimated from multiple shape descriptors. They aggregate this topological information from many quartets computed across the entire shape collection, and construct a hierarchical *categorization tree* (see Figure 20). Analogous to the phylogenetic trees of species, this categorization tree of a shape collection provides an overview of the shapes as well as their mutual distance and hierarchical relations. Based on such an organization, they also define the degree of separation chart for every shape in the collection and apply it for interactive shape exploration.

**Attribute-based exploration** Yet another approach seeks to allow users to interactively explore shapes with continuously valued semantic attributes. Blanz and Vetter [BV99] provide an interface to

explore faces based on continuous facial attributes, such as ‘smile’ or ‘frown’, built upon the face parametric model (Section 6). Similarly, Allen *et al.* [ACP03] allow users to explore the range of human bodies with features such as height, weight and age. Chaudhuri *et al.*’s [CKG\*13] interface enables exploration of shape parts according to learned strengths of semantic attributes (Figure 5).

## 10. Discussion

There is no ‘magic recipe’ for developing new data-driven and machine learning applications in geometry processing and computer graphics. Yet, there are some important considerations one needs to make in devising a data-driven method, including computational complexity, scalability, applicability issues, proper evaluation procedures and limitations. In this section, we briefly discuss these issues.

**Computational complexity.** As explained in Section 2, data-driven shape analysis and processing algorithms generally contain several stages (Figure 2). The complexity of each stage varies and largely depends on the number of input shapes, resolution of the input shape representation (number of faces, surface points, pixels or voxels), as well as the number and type of the used geometric features. The feature extraction stage is usually executed per each shape, thus, its time complexity often tends to be linear in the number of input shapes. Local geometric features, such as surface curvature or PCA-based descriptors, are usually computed within a small neighbourhood around each vertex, face or surface point, thus their extraction depends linearly on the number of these primitives in the input shape representation. Extracting geometric features that capture less local or global information about the shape, such as shape diameter, geodesic distance-based features or heat-kernel descriptors, is often computationally more intensive, i.e. super-linear in the number of primitives.

During the learning and inference steps, data-driven methods usually solve an optimization problem, which involves minimizing or maximizing a function, e.g. a data likelihood function. In general, optimization techniques inhabit a wide range of computational complexities. For example, if the optimization involves the least-squares solution of a linear system, as in the case of linear regression, the complexity is  $O(N \cdot F^2)$ , where  $N$  is the number of input training examples and  $F$  is the dimensionality of the input feature vector.

If optimization is performed through a steepest descent algorithm, the complexity is  $O(N \cdot F)$  per parameter update step. However, the performance of iterative optimization algorithms depends on the number of steps, which in turn varies according to their convergence properties and the function they optimize. We refer the reader to [NW06, KF09, Sol15] for an in-depth discussion on the computational complexity and convergence properties of various optimization and inference algorithms.

**Scalability.** Data-driven methods are inherently bound up with the input data. Rapid developments in capturing and modelling techniques have engendered the growth of 3D shape and scene repositories over recent years, which have in turn influenced the advancement of data-driven geometry processing. This is evidenced by the fact that the number of training shapes employed in data-driven geometry processing techniques has grown from a few tens [KHS10, SvKK\*11] to several thousands [KLM\*13, HWG14]. On the one hand, the increasing availability of 3D data can improve the accuracy and generalization of data-driven methods. On the other hand, issues of scalability arise. More data cause longer processing times, which in turn makes the debugging of such methods harder for developers. The scalability issues are further exacerbated by the complexity (i.e. high dimensionality) of the 3D geometric feature representations. Potential workarounds include debugging the pipeline of these methods on smaller data sets before turning to larger ones, trying simpler learning techniques before switching to more complex ones or making use of computing clusters for executing offline steps.

**Scope of application.** Not every problem in shape analysis and processing is well-suited to be solved by a data-driven method. When the underlying rules, principles and parameters can be manually and unambiguously specified in a problem, then non-data-driven methods should be considered for it. For example, deforming a shape with an elastic material and known physical parameters and forces can be addressed by a physics-based method rather than a data-driven one. In contrast, there are several problems in shape analysis and processing for which it is hard, or even impossible, to hand-design a set of rules and principles, or quantify them through manually specified parameters. This is often the case for problems that involve shape and scene recognition, high-level processing, structure parsing, co-analysis, reconstruction from noisy missing data, and modelling with high-level user input. Shape co-analysis (e.g. co-segmentation), in particular, requires estimating several possible geometric and semantic correlations among input shapes, which would be practically impossible to capture through hand-designed rules. Data-driven methods that automatically compute geometric, semantic and structural relationships in the input shapes are more appropriate for such co-analysis problems [XLZ\*10, HKG11, HG13]. Another example can be found in the problem of shape style analysis. Although humans have an innate sense of style similarity, it is hard to manually quantify geometric criteria for modelling the stylistic similarity of shapes. A style analysis algorithm whose parameters are learned through a data-driven method is much more well-suited to perform this quantification [LHLF15, LKS15].

**Evaluation.** Correctly evaluating the predictive performance of data-driven methods should be another important consideration for researchers of such methods. A common pitfall is to evaluate the

predictive performance of a data-driven method with the same data set on which it was trained on (e.g. through supervised learning), or a data set for which any parameters of the method were manually tuned. The risk here is that the method might not be able to generalize to any other data sets beyond the ones used in training or hand-tuning. A method that simply memorizes the training data set, or overfits a model to a particular data set, will obviously perform well there. However, if its performance on other data sets is poor, the method is effectively useless. The best practice is to introduce training and test splits of the input data sets. The models and parameters should then be learned or tuned exclusively on the training portion, and evaluated exclusively on the testing portion. To insure fair-play, it is also necessary that different data-driven methods be compared using the same training and test splits.

**Limitations** The data-driven approach to shape analysis and processing is bound by a few limitations that we summarize below. We also discuss potential workarounds to overcome some of these.

- **Generalization guarantees.** It is generally hard to provide any guarantees about the generalization performance of data-driven algorithms. In other words, when a data-driven algorithm makes use of a particular data set for training, it is often impossible to predict how well it will generalize to other data sets beforehand. Although statistical error bounds can be provided under particular assumptions on data distributions, in particular within the context of the Probably Approximately Correct learning theory [Val84] or the Bayes decision theory [Fuk90], these assumptions often cannot be validated in practice.
- **Complexity and scalability.** As discussed above, data-driven methods are computationally intensive in general. The complexity of data-driven methods depends on the number of input training shapes. As a general rule of thumb, the accuracy of data-driven methods improves with more training data. On the other hand, this comes at a higher computational cost during training time.
- **Size of 3D shape data sets.** Despite their recent growth, the size of available 3D shape data sets remains much smaller than those used in computer vision and natural language processing (e.g. image and text data sets). As a result, overfitting remains a common issue with data-driven methods for 3D shape processing. Overfitting occurs when a learned model, or function, captures random error, noise or patterns specific only to the input training data set instead of the underlying, correct relationships in the data. It usually occurs when the learned model, or function, is excessively complex, e.g. having an extremely large number of parameters relative to the size of the training data set. To mitigate this issue, regularization techniques can be used to favour simpler models and functions [Ng04, Dom12]. Another promising approach is to use both 3D shapes and 2D images as input to data-driven methods, or in other words, to perform co-analysis of image and shape data. We discuss this issue as one important future research direction for further development in data-driven methods in the next section.
- **Data collection.** Data-driven techniques rely on the availability of data for the particular problem they attempt to solve. Gathering training data for several geometry processing tasks is often not an easy task, especially when human labour is involved to process



or annotate geometric data. Although crowdsourcing Internet marketplaces, such as Amazon Mechanical Turk [Ama09], can help gather training data efficiently, online questionnaires and user studies still require careful design, monetary compensation and participant consistency checks.

- **From data to knowledge.** Data-driven methods put particular emphasis on discovering patterns and models that explain the input data and provide useful insights to the problem being solved. However, these learned patterns and models might not always be readily interpretable, i.e. might not correspond to ‘easy-to-understand’ rules. This is a common situation when one treats the internals of the data-driven method (e.g. the learning process) as a ‘black box’ without first trying to understand their exact functionality in detail. In general, interpreting such patterns and models requires significant time and effort.

## 11. Conclusion and Open Problems

In this survey, we have so far discussed state of the art on data-driven methods for 3D shape analysis and processing. We also presented the main concepts and methodologies used to develop such methods. We hope that this survey will act as a tutorial that will help researchers develop new data-driven algorithms related to shape analysis and processing. There are several exciting research directions that have not been sufficiently explored so far in our community that we discuss below:

**Joint analysis of 2D and 3D data.** Generating 3D content from images requires building mappings from 2D to 3D space. Unfortunately, the problem remains largely ill-posed. However, with the help vast quantities of 2D images available on the web, effective priors can be developed to map 2D visual elements or features to 3D shape and scene representations. Indeed, we have in fact seen recent attempts made in this very vein of thought with some success in [SHM\*14, AME\*14, LSQ\*15, HOM15, SQLG15], which attempts depth estimation through joint analysis over 2D image collections and 3D model databases. We have also seen success of the joint analysis framework in the setting of texture-data with [YCM14], which attempts cosegmentation of textured 3D shapes.

Following this line, it would be interesting to jointly analyse and process multi-modal visual data, including depth scans and videos. The key challenge lies in the integration of heterogeneous information in a unified learning framework.

**Better and scalable shape analysis techniques.** Many data-driven applications rely on high-quality shape analysis results, particularly shape segmentations and correspondences. We believe it is important to further advance research in both these directions. This includes designing shape analysis techniques for specific data and/or making them scalable to very large data sets, especially recently emerging large-scale richly annotated repositories [SSY\*15].

**From geometry to semantics and vice versa.** Several data-driven methods have tried to map 2D and 3D geometric data to high-level concepts, such as shape categories, semantic attributes or part labels. Gathering relevant training data is a key component in achieving this aim, a task which remains a non-trivial endeavour.

Several recent promising works employ crowdsourcing to address this issue [CGF09, CKG\*13, LKS15, LHLF15, YCHK15]. Existing methods deal with cases where only a handful of different entities are predicted for input shapes or scenes. Scaling these methods to handle thousands of categories, part labels and other such entities, as well as attaining human-level performance, is an open problem. The opposite direction is also interesting and insufficiently explored: generating shapes and scenes based on high-level specifications such as shape styles, attributes or even natural language. Such approaches may even potentially be combined with further diverse inputs, such as sketches and interactive handles, in the shape-generating pipeline. WordsEye [CS01] was an early attempt to bridge this gap, yet requires extensive manual mappings.

**Understanding function from geometry.** The geometry of a shape is strongly related to its functionality, including the shape’s relationship to human activity. Thus, analysing shapes and scenes requires some understanding of their function. The recent works by Laga *et al.* [LMS13], Kim *et al.* [KCGF14] and Hu *et al.* [HZvK\*15] are important examples of data-driven approaches that take into account functional aspects of shapes in the process of their analysis. In addition, data-driven methods can guide the synthesis of shapes that can be manufactured or 3D printed based on given functional specifications; an example of such an attempt is reflected in the work of Schulz *et al.* [SSL\*14].

**Data-driven shape abstractions.** It is relatively easy for humans to communicate the essence of shapes with a few lines, sketches and abstract forms. Developing methods that can build such abstractions automatically has significant applications in shape and scene visualization, artistic rendering and shape analysis. There are a few data-driven approaches to line drawing [CGL\*08, KNS\*09, KNBH12], saliency analysis [CSPF12], surface abstraction [YK12] and viewpoint preferences [SLF\*11] related to this goal. Matching human performance in these tasks is still a largely open problem, while synthesizing and editing shapes using shape abstractions as input remains a significant challenge.

**Feature learning.** Several shape and scene processing tasks depend on engineering geometric features for points and shapes, as discussed in Section 3. In general, it seems that some features work well in particular settings, but can fail in others. A prevailing issue is that universal geometric descriptors—features that can serve as reliable mid or high-level representations ubiquitously across all variety of shapes—do not yet exist.

Recent work in machine learning has demonstrated that powerful feature representations can be learned directly from raw input text and image data with deep architectures [HOT06, KSH12, ZF14]. These architectures are composed of multiple processing layers which learn representations of the input data at multiple levels of abstraction. These data-driven representations are optimized for processing-performance in complex interpretation tasks. Such feature learning for 3D shapes with deep architectures has recently been demonstrated in the context of shape classification [WSK\*15, SMKLM15, XXS\*15, HKM15]. Learning features for performing other complex high-level shape analysis and processing tasks remains an open problem.

## Acknowledgements

We thank Zimo Li for proofreading this survey and the anonymous reviewers for helpful suggestions. Kalogerakis gratefully acknowledges support from NSF (CHS-1422441). Kai Xu is supported by NSFC (61572507, 61202333 and 61532003).

## References

- [ACP03] ALLEN B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: Reconstruction and parameterization from range scans. *ACM Transactions on Graphics* 22, 3 (2003), 587–594.
- [AKKS99] ANKERST M., KASTENMÜLLER G., KRIEGEL H.-P., SEIDL T.: 3D shape histograms for similarity search and classification in spatial databases. In *Proceedings of SSD '99* (Springer, London, UK, 1999), pp. 207–226.
- [AKZM14] AVERKIOU M., KIM V. G., ZHENG Y., MITRA N. J.: ShapeSynth: Parameterizing model collections for coupled shape exploration and synthesis. *Computer Graphics Forum* 33, 2 (2014), 125–134.
- [Ama09] Amazon: Mechanical Turk. <http://www.mturk.com> (2009).
- [AME\*14] AUBRY M., MATURANA D., EFROS A. A., RUSSELL B. C., SIVIC J.: Seeing 3D chairs: Exemplar part-based 2D-3D alignment using a large dataset of cad models. In *Proceedings of CVPR* (Columbus, OH, USA, 2014).
- [ASK\*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: Shape completion and animation of people. In *Proceedings of SIGGRAPH* (New York, NY, USA, 2005), pp. 408–416.
- [ATC\*05] ANGUELOV D., TASKAR B., CHATALBASHEV V., KOLLER D., GUPTA D., HEITZ G., NG A.: Discriminative learning of Markov random fields for segmentation of 3D scan data. In *Proceedings of CVPR* (Washington, DC, USA, 2005).
- [BB01] BANKO M., BRILL E.: Mitigating the paucity-of-data problem: Exploring the effect of training corpus size on classifier performance for natural language processing. In *Proceedings of International Conference on Human Language Technology Research* (Stroudsburg, PA, USA, 2001), pp. 1–5.
- [BBOG11] BRONSTEIN A. M., BRONSTEIN M. M., OVSJANIKOV M., GUIBAS L. J.: Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics* 30, 1 (2011), 1–20.
- [BD06] BARUTCUOGLU Z., DECORO C.: Hierarchical shape classification using Bayesian aggregation. In *Proceedings of Shape Modeling International* (Matsushima, Japan, 2006).
- [Ben09] BENGIO Y.: Learning deep architectures for AI. *Foundations and Trends in Machine Learning* 2, 1 (2009), 1–127.
- [BLVD11] BENHABILES H., LAVOUÉ G., VANDEBORRE J.-P., DAUDI M.: Learning boundary edges for 3D-mesh segmentation. *Computer Graphics Forum* 30, 8 (2011), 2170–2182.
- [BMM\*15] BOSCAINI D., MASCI J., MELZI S., BRONSTEIN M. M., CASTELLANI U., VANDERGHEYNST P.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Computer Graphics Forum* 34, 5 (2015), 13–23.
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 4 (2002), 509–522.
- [BRF14] BO L., REN X., FOX D.: Learning hierarchical sparse features for RGB-(D) object recognition. *International Journal of Robotics Research* 33, 4, (2014), 581–599.
- [BSWR12] BLUM M., SPRINGENBERG J. T., WULFING J., RIEDMILLER M.: A learned feature descriptor for object recognition in RGB-D data. In *Proceedings of IEEE International Conference on Robotics and Automation* (2012), pp. 1298–1303.
- [BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *Proceedings of SIGGRAPH* (New York, NY, USA, 1999), pp. 187–194.
- [CAF10a] CHO T. S., AVIDAN S., FREEMAN W. T.: The patch transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 8 (2010), 1489–1501.
- [CAF10b] CHO T. S., AVIDAN S., FREEMAN W. T.: A probabilistic image jigsaw puzzle solver. In *Proceedings of CVPR* (San Francisco, CA, USA, 2010), pp. 183–190.
- [CAK12] CAMPEN M., ATTENE M., KOBELT L.: A practical guide to polygon mesh repairing. In *Eurographics Tutorials* (2012).
- [CCPS13] CHOI W., CHAO Y.-W., PANTOFARU C., SAVARESE S.: Understanding indoor scenes using 3D geometric phrases. In *Proceedings of CVPR* (Portland, OR, USA, 2013).
- [CGF09] CHEN X., GOLOVINSKIY A., FUNKHOUSER T.: A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics* 28, 3 (2009), 73:1–73:12.
- [CGH14] CHEN Y., GUIBAS L. J., HUANG Q.-X.: Near-optimal joint object matching via convex relaxation. CoRR abs/1402.1473 (2014).
- [CGL\*08] COLE F., GOLOVINSKIY A., LIMPAECHER A., BARROS H. S., FINKELSTEIN A., FUNKHOUSER T., RUSINKIEWICZ S.: Where do people draw lines? *ACM Transactions on Graphics* 27, 3 (2008), 88:1–88:11.
- [CHZ14] CAO C., HOU Q., ZHOU K.: Displaced dynamic expression regression for real-time facial tracking and animation. *ACM Transactions on Graphics* 33, 4 (2014), 43:1–43:10.
- [CK10] CHAUDHURI S., KOLTUN V.: Data-driven suggestions for creativity support in 3D modeling. *ACM Transactions on Graphics* 29, 6 (2010), 183:1–183:10.
- [CKG\*13] CHAUDHURI S., KALOGERAKIS E., GIGUERE S., FUNKHOUSER T.: *AttribIt*: Content creation with semantic

- attributes. In *Proceedings of UIST* (New York, NY, USA, 2013).
- [CKGK11] CHAUDHURI S., KALOGERAKIS E., GUIBAS L., KOLTUN V.: Probabilistic reasoning for assembly-based 3D modeling. *ACM Transactions on Graphics* 30, 4 (2011), 35:1–35:10.
- [CLM\*12] CHANG W., LI H., MITRA N., PAULY M., WAND M.: Dynamic geometry processing. In *Eurographics Tutorials* (2012).
- [CLW\*14] CHEN K., LAI Y.-K., WU Y.-X., MARTIN R., HU S.-M.: Automatic semantic modeling of indoor scenes from low-quality RGB-D data using contextual information. *ACM Transactions on Graphics* 33, 6 (2014), 208:1–208:12.
- [COSH11] CRANDALL D., OWENS A., SNAVELY N., HUTTENLOCHER D.: Discrete-continuous optimization for large-scale structure from motion. In *Proceedings of CVPR '11* (Washington, DC, USA, 2011), pp. 3001–3008.
- [CS01] COYNE B., SPROAT R.: Wordseye: An automatic text-to-scene conversion system. In *Proceedings of SIGGRAPH* (New York, NY, USA, 2001).
- [CSPF12] CHEN X., SAPAROV A., PANG B., FUNKHOUSER T.: Schelling points on 3D surface meshes. *ACM Transactions on Graphics* 31, 4 (2012), 29:1–29:12.
- [CTS03] CHEN D.-Y., TIAN X.-P., SHEN Y.-T., OUHYOUNG M.: On visual similarity based 3D model retrieval. *Computer Graphics Forum* 22, 3 (2003), 223–232.
- [CWLZ13] CAO C., WENG Y., LIN S., ZHOU K.: 3D shape regression for real-time facial animation. *ACM Transactions on Graphics* 32, 4 (2013), 41:1–41:10.
- [CWZ\*14] CHEN C., WENG Y., ZHOU S., TONG Y., ZHOU K.: Face-warehouse: A 3D facial expression database for visual computing. *IEEE Transactions on Visualization & Computer Graphics* 20, 3 (2014), 413–425.
- [Dom12] DOMINGOS P.: A few useful things to know about machine learning. *Communications of the ACM* 55, 10 (2012), 78–87.
- [DTB06] DIEBEL J., THRUN S., BRUENING M.: A Bayesian method for probable surface reconstruction and decimation. *ACM Transactions on Graphics* 25, 1 (2006), 39–59.
- [FAvK\*14] FISH N., AVERKIOU M., VAN KAICK O., SORKINE-HORNUNG O., COHEN-OR D., MITRA N. J.: Meta-representation of shape families. *ACM Transactions on Graphics* 33, 4 (2014), 34:1–34:11.
- [FFFP06] FEI-FEI L., FERGUS R., PERONA P.: One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 4 (2006), 594–611.
- [FGH13] FOUHEY D. F., GUPTA A., HEBERT M.: Data-driven 3D primitives for single image understanding. In *Proceedings of ICCV* (Sydney, NSW, 2013).
- [FH10] FISHER M., HANRAHAN P.: Context-based search for 3D models. *ACM Transactions on Graphics* 29, 6 (2010), 182:1–182:10.
- [FHK\*04] FROME A., HUBER D., KOLLURI R., BULOW T., MALIK J.: Recognizing objects in range data using regional point descriptors. In *Proceedings of ECCV* (Prague, Czech Republic, 2004), pp. 224–237.
- [FKMS05] FUNKHOUSER T. A., KAZHDAN M. M., MIN P., SHILANE P.: Shape-based retrieval and analysis of 3D models. *Communications of the ACM* 48, 6 (2005), 58–64.
- [FKS\*04] FUNKHOUSER T., KAZHDAN M., SHILANE P., MIN P., KIEFER W., TAL A., RUSINKIEWICZ S., DOBKIN D.: Modeling by example. *ACM Transactions on Graphics* 23, 3 (2004), 652–663.
- [FRS\*12] FISHER M., RITCHIE D., SAVVA M., FUNKHOUSER T., HANRAHAN P.: Example-based synthesis of 3D object arrangements. *ACM Transactions on Graphics* 31, 6 (2012), 135:1–135:12.
- [FS06] FUNKHOUSER T., SHILANE P.: Partial matching of 3D shapes with priority-driven search. In *Proceedings of Symposium on Geometry Processing* (Cagliari, Sardinia, 2006).
- [FSH11] FISHER M., SAVVA M., HANRAHAN P.: Characterizing structural relationships in scenes using graph kernels. *ACM Transactions on Graphics* 30, 4 (2011), 34:1–34:12.
- [Fuk90] FUKUNAGA K.: *Introduction to Statistical Pattern Recognition*. Academic Press Professional, Inc., 1990, San Diego, CA, USA.
- [FWX\*13] FAN L., WANG R., XU L., DENG J., LIU L.: Modeling by drawing with shadow guidance. *Computer Graphics Forum* 32, 7 (2013), 157–166.
- [GF09] GOLOVINSKIY A., FUNKHOUSER T. A.: Consistent segmentation of 3D models. *Computers & Graphics* 33, 3 (2009), 262–269.
- [GKF09] GOLOVINSKIY A., KIM V. G., FUNKHOUSER T.: Shape-based recognition of 3D point clouds in urban environments. In *Proceedings of ICCV* (Kyoto, Japan, 2009), pp. 2154–2161.
- [GLXJ14] GUO X., LIN J., XU K., JIN X.: Creature grammar for creative modeling of 3D monsters. *Graphical Models (Special Issue of GMP)* 76, 5 (2014), 376–389.
- [GMB04] GOLDBERG D., MALON C., BERN M.: A global approach to automatic solution of jigsaw puzzles. *Computational Geometry: Theory and Applications* 28 (2004), 165–174.
- [GSH\*07] GAL R., SHAMIR A., HASSNER T., PAULY M., COHEN-OR D.: Surface reconstruction using local shape priors. In *SGP '07: Proceedings of Symposium on Geometry Processing* (Barcelona, Spain, 2007), pp. 253–262.
- [GSMCO09] GAL R., SORKINE O., MITRA N. J., COHEN-OR D.: iwires: An analyze-and-edit approach to shape manipulation. *ACM Transactions on Graphics* (2009), 33:1–33:10.

- [HCG14] HUANG Q., CHEN Y., GUIBAS L. J.: Scalable semidefinite relaxation for maximum a posterior estimation. In *Proceedings of International Conference on Machine Learning (ICML)* (Beijing, China, 2014).
- [HFG\*06] HUANG Q., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. *ACM Transactions on Graphics* 25, 3 (2006), 569–578.
- [HFL12] HU R., FAN L., LIU L.: Co-segmentation of 3D shapes via subspace clustering. *Computer Graphics Forum* 31, 5 (2012), 1703–1713.
- [HG13] HUANG Q., GUIBAS L.: Consistent shape maps via semidefinite programming. *Computer Graphics Forum* 32, 5 (2013), 177–186.
- [HKG11] HUANG Q., KOLTUN V., GUIBAS L.: Joint shape segmentation using linear programming. *ACM Transactions on Graphics* 30, 6 (2011), 125:1–125:12.
- [HKM15] HUANG H., KALOGERAKIS E., MARLIN B.: Analysis and synthesis of 3D shape families via deep-learned generative models of surfaces. *Computer Graphics Forum* 34, 5 (2015), 25–38.
- [HOM15] HUETING M., OVSIANIKOV M., MITRA N. J.: Crosslink: Joint understanding of image and 3D model collections through shape and camera pose variations. *ACM Transactions on Graphics* (2015), 34, 6, 233:1–233:13.
- [Hor84] HORN B. K. P.: Extended Gaussian images. *Proceedings of the IEEE* 72, 12 (1984), 1671–1686.
- [HOT06] HINTON G. E., OSINDERO S., TEH Y.-W.: A fast learning algorithm for deep belief nets. *Neural Computation* 18, 7 (2006), 1527–1554.
- [HSG13] HUANG Q., SU H., GUIBAS L.: Fine-grained semi-supervised labeling of large shape collections. *ACM Transactions on Graphics* 32, 6 (2013), 190:1–190:10.
- [HSR\*09] HASLER N., STOLL C., ROSENHAHN B., THORMÄHLEN T., SEIDEL H.-P.: Technical section: Estimating body shape of dressed humans. *Computers & Graphics* 33, 3 (2009), 211–216.
- [HSS\*09] HASLER N., STOLL C., SUNKEL M., ROSENHAHN B., SEIDEL H.-P.: A statistical model of human pose and body shape. *Computer Graphics Forum* 28, 2 (2009), 337–346.
- [HSS\*13] HUANG S.-S., SHAMIR A., SHEN C.-H., ZHANG H., SHEFFER A., HU S.-M., COHEN-OR D.: Qualitative organization of collections of shapes via quartet analysis. *ACM Transactions on Graphics* 32, 4 (2013), 71:1–71:10.
- [Hub02] HUBER D.: Automatic Three-Dimensional Modeling from Reality. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 2002.
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics* 33, 4 (2014), 36:1–36:11.
- [HWK15] HUANG Q., WANG H., KOLTUN V.: Single-view reconstruction via joint analysis of image and shape collections. *ACM Transactions on Graphics* 34 (2015), 87:1–87:10.
- [HZG\*12] HUANG Q., ZHANG G.-X., GAO L., HU S.-M., BUTSCHER A., GUIBAS L.: An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Transactions on Graphics* 31, 6 (2012), 167:1–167:11.
- [HZvK\*15] HU R., ZHU C., VAN KAICK O., LIU L., SHAMIR A., ZHANG H.: Interaction context (icon): Towards a geometric functionality descriptor. *ACM Transactions on Graphics* 34, 4 (2015), 83:1–83:12.
- [JH99] JOHNSON A. E., HEBERT M.: Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 5 (1999), 433–449.
- [JWL\*06] JIAO F., WANG S., LEE C.-H., GREINER R., SCHUURMANS D.: Semi-supervised conditional random fields for improved sequence segmentation and labeling. In *Proceedings of 21st International Conference on Computational Linguistics* (Stroudsburg, PA, USA, 2006).
- [KBLB12] KOKKINOS I., BRONSTEIN M., LITMAN R., BRONSTEIN A.: Intrinsic shape context descriptors for deformable shapes. In *Proceedings of CVPR* (2012, Providence, RI, USA).
- [KCGF14] KIM V. G., CHAUDHURI S., GUIBAS L., FUNKHOUSER T.: Shape2Pose: Human-centric shape analysis. *ACM Transactions on Graphics* 33, 4 (2014), 120:1–120:12.
- [KCKK12] KALOGERAKIS E., CHAUDHURI S., KOLLER D., KOLTUN V.: A probabilistic model for component-based shape synthesis. *ACM Transactions on Graphics* 31, 4 (2012), 55:1–55:11.
- [KF09] KOLLER D., FRIEDMAN N.: *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009, Cambridge, MA, USA.
- [KFR03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Proceedings of Symposium on Geometry Processing* (2003).
- [KFR04] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Symmetry descriptors and 3D shape matching. In *Proceedings of Symposium on Geometry Processing* (Aachen, Germany, 2004).
- [KHS10] KALOGERAKIS E., HERTZMANN A., SINGH K.: Learning 3D mesh segmentation and labeling. *ACM Transactions on Graphics* 29 (2010), 102:1–102:12.
- [KJS07] KREAVOY V., JULIUS D., SHEFFER A.: Model composition from interchangeable components. In *Proceedings of Pacific Graphics* (Maui, Hawaii, USA, 2007), IEEE Computer Society, pp. 129–138.
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Transactions on Graphics* 30, 4 (2011), 79:1–79:12.

- [KLM\*12] KIM V. G., LI W., MITRA N. J., DiVERDI S., FUNKHOUSER T.: Exploring collections of 3D models using fuzzy correspondences. *ACM Transactions on Graphics* 31, 4 (2012), 54:1–54:11.
- [KLM\*13] KIM V. G., LI W., MITRA N. J., CHAUDHURI S., DiVERDI S., FUNKHOUSER T.: Learning part-based templates from large collections of 3D shapes. *ACM Transactions on Graphics* 32, 4 (2013), 70:1–70:12.
- [KMYG12] KIM Y. M., MITRA N. J., YAN D.-M., GUIBAS L.: Acquiring 3D indoor environments with variability and repetition. *ACM Transactions on Graphics* 31, 6 (2012), 138:1–138:11.
- [KNBH12] KALOGERAKIS E., NOWROUZEZAHRAI D., BRESLAV S., HERTZMANN A.: Learning hatching for pen-and-ink illustration of surfaces. *ACM Transactions on Graphics* 31, 1 (2012), 1:1–1:17.
- [KNS\*09] KALOGERAKIS E., NOWROUZEZAHRAI D., SIMARI P., McCRAE J., HERTZMANN A., SINGH K.: Data-driven curvature for real-time line drawing of dynamic scenes. *ACM Transactions on Graphics* 28, 1 (2009), 1–13.
- [KSES14] KHOLOGADE N., SIMON T., EFROS A., SHEIKH Y.: 3D object manipulation in a single photograph using stock 3D models. *ACM Transactions on Graphics* 33, 4 (2014), 127:1–127:12.
- [KSH12] KRIZHEVSKY A., SUTSKEVER I., HINTON G. E.: Imagenet classification with deep convolutional neural networks. In *Proceedings of NIPS* (Lake Tahoe, NV, USA, 2012).
- [KSNS07] KALOGERAKIS E., SIMARI P., NOWROUZEZAHRAI D., SINGH K.: Robust statistical estimation of curvature on discretized surfaces. In *Proceedings of Symposium on Geometry Processing* (Barcelona, Spain, 2007).
- [KT03] KATZ S., TAL A.: Hierarchical mesh decomposition using fuzzy clustering and cuts. In *Proceedings of ACM SIGGRAPH '03* (San Diego, CA, USA), pp. 954–961.
- [KWT88] KASS M., WITKIN A., TERZOPOULOS D.: Snakes: Active contour models. *International Journal of Computer Vision* 1, 4 (1988), 321–331.
- [LBBC14] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U.: Supervised learning of bag-of-features shape descriptors using sparse coding. In *Proceedings of SGP* (Cardiff, UK, 2014).
- [LBF13] LAI K., BO L., FOX D.: Unsupervised feature learning for 3D scene labeling. In *Proceedings of IEEE International Conference on Robotics and Automation* (Karlsruhe, Germany, 2013).
- [LCHB12] LV J., CHEN X., HUANG J., BAO H.: Semi-supervised mesh segmentation and labeling. *Computer Graphics Forum* 31, 7-2 (2012), 2241–2248.
- [LCK\*14] LIU T., CHAUDHURI S., KIM V., HUANG Q.-X., MITRA N. J., FUNKHOUSER T.: Creating consistent scene graphs using a probabilistic grammar. *ACM Transactions on Graphics* 33, 6 (2014), 211:1–211:12.
- [LDGN15] LI Y., DAI A., GUIBAS L., NIESSNER M.: Database-assisted object retrieval for real-time 3D reconstruction. *Computer Graphics Forum* 34 (2015), 435–446.
- [LF09] LIPMAN Y., FUNKHOUSER T.: Mobius voting for surface correspondence. *ACM Transactions on Graphics* 28, 3 (2009), 72:1–72:12.
- [LGA\*12] LI B., GODIL A., AONO M., BAI X., FURUYA T., LI L., LOPEZ-SASTRE R., JOHAN H., OHBUCHI R., REDONDO-CABRERA C., TATSUMA A., YANAGIMACHI T., ZHANG S.: SHREC' 12 track: Generic 3D shape retrieval. In *Proceedings of Eurographics Workshop on 3D Object Retrieval* (Cagliari, Italy, 2012).
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proceedings of ICCV '05* (Beijing, China), pp. 1482–1489.
- [LHLF15] LIU T., HERTZMANN A., LI W., FUNKHOUSER T.: Style compatibility for 3D furniture models. *ACM Transactions on Graphics* 34, 4 (2015), 85:1–85:9.
- [LKS15] LUN Z., KALOGERAKIS E., SHEFFER A.: Elements of style: Learning perceptual shape style similarity. *ACM Transactions on Graphics* 34, 4 (2015), 84:1–84:14.
- [LLL\*14] LI B., LU Y., LI C., GODIL A., SCHRECK T., AONO M., CHEN Q., CHOWDHURY N. K., FANG B., FURUYA T., JOHAN H., KOSAKA R., KOYANAGI H., OHBUCHI R., TATSUMA A.: SHREC' 14 track: Large scale comprehensive 3D shape retrieval. In *Proceedings of Eurographics Workshop on 3D Object Retrieval* (Strasbourg, France, 2014), pp. 131–140.
- [LMB14] LOPER M. M., MAHMOOD N., BLACK M. J.: MoSh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics* 33, 6 (2014), 220:1–220:13.
- [LMP01] LAFFERTY J. D., MCCALLUM A., PEREIRA F. C. N.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML* (2001), pp. 282–289.
- [LMS13] LAGA H., MORTARA M., SPAGNUOLO M.: Geometry and context for semantic correspondences and functionality recognition in man-made 3D shapes. *ACM Transactions on Graphics* 32, 5 (2013), 150:1–150:16.
- [LSQ\*15] LI Y., SU H., QI C. R., FISH N., COHEN-OR D., GUIBAS L.: Joint embeddings of shapes and images via CNN image purification. *ACM Transactions on Graphics* (2015), 34, 6, 234:1–234:12.
- [LWP10] LI H., WEISE T., PAULY M.: Example-based facial rigging. *ACM Transactions on Graphics* 29, 4 (2010), 32:1–32:6.
- [LYYB13] LI H., YU J., YE Y., BREGLER C.: Realtime facial animation with on-the-fly correctives. *ACM Transactions on Graphics* 32, 4 (2013), 42:1–42:10.
- [LZ11] LEVY B., ZHANG H.: Elements of geometry processing. In *Invited SIGGRAPH Asia Courses* (2011).

- [LZC11] LEE Y. J., ZITNICK L., COHEN M.: Shadowdraw: Real-time user guidance for freehand drawing. *ACM Transactions on Graphics* 30, 4 (2011), 27:1–27:9.
- [Mer07] MERRELL P.: Example-based model synthesis. In *Proceedings of I3D* (Seattle, WA, USA, 2007), pp. 105–112.
- [MHS\*14] MA C., HUANG H., SHEFFER A., KALOGERAKIS E., WANG R.: Analogy-driven 3D style transfer. *Computer Graphics Forum* 33, 2 (2014), 175–184.
- [MPM\*14] MATTAUSCH O., PANOZZO D., MURA C., SORKINE-HORNUNG O., PAJAROLA R.: Object detection and classification from large-scale cluttered indoor scans. *Computer Graphics Forum* 33, 2 (2014), 11–21.
- [MRS08] MANNING C. D., RAGHAVAN P., SCHÜTZE H.: *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.
- [MWH\*06] MÜLLER P., WONKA P., HAEGLER S., ULMER A., VAN GOOL L.: Procedural modeling of buildings. *ACM Transactions on Graphics* (2006), 614–623.
- [MWZ\*14] MITRA N., WAND M., ZHANG H., COHEN-OR D., KIM V., HUANG Q.-X.: Structure-aware shape processing. In *ACM SIG-GRAPH Courses* (2014).
- [NBCW\*11] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. *Computer Graphics Forum* 30, 5 (2011), 1481–1491.
- [Ng04] NG A. Y.: Feature selection,  $l_1$  vs.  $l_2$  regularization, and rotational invariance. In *Proceedings of ICML* (Banff, Canada, 2004).
- [NK03] NOVOTNI M., KLEIN R.: 3D zernike descriptors for content based shape retrieval. In *Proceedings of Solid Modeling* (Seattle, WA, USA, 2003).
- [NW06] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization*. Springer, New York, 2006.
- [NXS12] NAN L., XIE K., SHARF A.: A search-classify approach for cluttered indoor scene understanding. *ACM Transactions on Graphics* 31, 6 (2012), 137:1–137:10.
- [OBCS\*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics* 31, 4 (2012), 30:1–30:11.
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Transactions on Graphics* 21 (2002), 807–832.
- [OLGM11] OVSJANIKOV M., LI W., GUIBAS L., MITRA N. J.: Exploration of continuous variability in collections of 3D shapes. *ACM Transactions on Graphics* 30, 4 (2011), 33:1–33:10.
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L. J.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 5 (2010), 1555–1564.
- [PEP\*11] PAULOVICH F., ELER D., POCO J., BOTHA C., MINGHIM R., NONATO L.: Piecewise Laplacian-based projection for interactive data exploration and organization. *Computer Graphics Forum* 30, 3 (2011), 1091–1100.
- [PKS13] PACHAURI D., KONDOR R., SINGH V.: Solving the multi-way matching problem by permutation synchronization. In *Proceedings of NIPS* (Lake Tahoe, NV, USA, 2013), pp. 1860–1868.
- [PMG\*05] PAULY M., MITRA N. J., GIESEN J., GROSS M., GUIBAS L. J.: Example-based 3D scan completion. In *SGP '05: Proceedings of Symposium on Geometry Processing* (Vienna, Austria, 2005).
- [QT09] QUATTONI A., TORRALBA A.: Recognizing indoor scenes. In *Proceedings of CVPR* (Miami, FL, USA, 2009), pp. 413–420.
- [ROA\*13] RUSTAMOV R. M., OVSJANIKOV M., AZENCOT O., BEN-CHEN M., CHAZAL F., GUIBAS L.: Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics* 32, 4 (2013), 72:1–72:12.
- [RR13] REN X., RAMANAN D.: Histograms of sparse codes for object detection. In *Proceedings of CVPR* (Portland, OR, USA, 2013), pp. 3246–3253.
- [RSSS11] ROBERTS R., SINHA S. N., SZELISKI R., STEEDLY D.: Structure from motion for scenes with large duplicate structures. In *Proceedings of CVPR* (Colorado Springs, CO, USA, 2011), pp. 3137–3144.
- [SF07] SHILANE P., FUNKHOUSER T.: Distinctive regions of 3D surfaces. *ACM Transactions on Graphics* 26, 2 (2007), pp. 7:1–7:15.
- [SFC\*11] SHOTTON J., FITZGIBBON A., COOK M., SHARP T., FINOCCHIO M., MOORE R., KIPMAN A., BLAKE A.: Real-time human pose recognition in parts from single depth images. In *Proceedings of CVPR* (Colorado Springs, CO, USA, 2011).
- [SFCH12] SHEN C.-H., FU H., CHEN K., HU S.-M.: Structure recovery by part assembly. *ACM Transactions on Graphics* 31, 6 (2012), 180:1–180:11.
- [Sha08] SHAMIR A.: A survey on mesh segmentation techniques. *Computer Graphics Forum* 27, 6 (2008), 1539–1556.
- [SHB\*12] SOCHER R., HUVAL B., BHAT B., MANNING C. D., NG A. Y.: Convolutional-recursive deep learning for 3D object classification. In *Proceedings of NIPS* (Lake Tahoe, NV, USA, 2012), pp. 665–673.
- [SHM\*14] SU H., HUANG Q.-X., MITRA N. J., LI Y., GUIBAS L. J.: Estimating image depth using shape collections. *ACM Transactions on Graphics* 33, 4 (2014), 37:1–37:11.
- [SKAG15] SUNG M., KIM V. G., ANGST R., GUIBAS L.: Data-driven structural priors for shape completion. *ACM Transactions on Graphics* 34, 6 (2015), 175:1–175:11.
- [SKHF12] SILBERMAN N., KOHLI P., HOIEM D., FERGUS R.: Indoor segmentation and support inference from RGBD images. In *Proceedings of ECCV* (Florence, Italy, 2012).

- [SLF\*11] SECORD A., LU J., FINKELSTEIN A., SINGH M., NEALEN A.: Perceptual models of viewpoint preference. *ACM Transactions on Graphics* 30, 5 (2011), 109:1–109:13.
- [SLH12] SATKIN S., LIN J., HEBERT M.: Data-driven scene understanding from 3D models. In *Proceedings of BMVC* (Guildford, UK, 2012), pp. 128:1–128:11.
- [SLJ\*15] SZEGEDY C., LIU W., JIA Y., SERMANET P., REED S., ANGUELOV D., ERHAN D., VANHOUCHE V., RABINOVICH A.: Going deeper with convolutions. In *Proceedings of CVPR* (Boston, MA, USA, 2015).
- [SM05] SUTTON C., MCCALLUM A.: Piecewise training of undirected models. In *Proceedings of UAI* (Edinburgh, Scotland, 2005).
- [SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The Princeton shape benchmark. In *Proceedings of Shape Modeling International* (Genova, Italy, 2004).
- [SMKLM15] SU H., MAJI S., KALOGERAKIS E., LEARNED-MILLER E.: Multi-view convolutional neural networks for 3D shape recognition. In *Proceedings of ICCV* (Santiago, Chile, 2015).
- [SMNS\*13] SALAS-MORENO R. F., NEWCOMBE R. A., STRASDAT H., KELLY P. H. J., DAVISON A. J.: SLAM++: Simultaneous localisation and mapping at the level of objects. In *Proceedings of CVPR* (Portland, OR, USA, 2013), IEEE, pp. 1352–1359.
- [Sol15] SOLOMON J.: *Numerical Algorithms*. AK Peters/CRC Press, Natick, MA, 2015.
- [SQLG15] SU H., QI C. R., LI Y., GUIBAS L. J.: Render for CNN: Viewpoint estimation in images using CNNs trained with rendered 3D model views. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (Santiago, Chile, December 2015).
- [SSCO08] SHAPIRA L., SHAMIR A., COHEN-OR D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. *Visual Computer* 24 (2008), 249–259.
- [SSL\*14] SCHULZ A., SHAMIR A., LEVIN D. I. W., SITTHI-AMORN P., MATUSIK W.: Design and fabrication by example. *ACM Transactions on Graphics* 33, 4 (2014), 62:1–62:11.
- [SSS\*10] SHAPIRA L., SHALOM S., SHAMIR A., ZHANG R. H., COHEN-OR D.: Contextual part analogies in 3D objects. *International Journal of Computer Vision* 89, 2 (2010), 309–315.
- [SSY\*15] SU H., SAVVA M., YI L., CHANG A. X., SONG S., YU F., LI Z., XIAO J., HUANG Q., SAVARESE S., FUNKHOUSER T., HANRAHAN P., GUIBAS L. J.: ShapeNet: An information-rich 3D model repository. <http://shapenet.cs.stanford.edu>.
- [SV01] SAUPE D., VRANIC D. V.: 3D model retrieval with spherical harmonics and moments. In *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition* (Munich, Germany, 2001), pp. 392–397.
- [SVKK\*11] SIDI O., VAN KAICK O., KLEIMAN Y., ZHANG H., COHEN-OR D.: Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Transactions on Graphics* 30, 6 (2011), 126:1–126:10.
- [SW10] SWADZBA A., WACHSMUTH S.: Indoor scene classification using combined 3D and gist features. In *Proceedings of ACCV* (Queenstown, New Zealand, 2010), pp. 201–215.
- [SW11] SINGER A., WU H.-T.: Vector diffusion maps and the connection Laplacian. *Communications on Pure and Applied Mathematics* 65, 8 (2011), 1067–1144.
- [SX14] SONG S., XIAO J.: Sliding shapes for 3D object detection in depth images. In *Proceedings of ECCV* (Zurich, Switzerland, 2014).
- [SXZ\*12] SHAO T., XU W., ZHOU K., WANG J., LI D., GUO B.: An interactive approach to semantic modeling of indoor scenes with an RGBD camera. *ACM Transactions on Graphics* 31, 6 (2012), 136:1–136:11.
- [SY07] SCHAEFER S., YUKSEL C.: Example-based skeleton extraction. In *Proceedings of Symposium on Geometric Processing* (Barcelona, Spain, 2007), pp. 153–162.
- [TFF08] TORRALBA A., FERGUS R., FREEMAN W. T.: 80 million tiny images: A large database for non-parametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, 11 (2008), 1958–1970.
- [TGY\*09] TALTON J. O., GIBSON D., YANG L., HANRAHAN P., KOLTUN V.: Exploratory modeling with collaborative design spaces. *ACM Transactions on Graphics* 28, 5 (2009), 167:1–167:10.
- [TMB14] TSOLI A., MAHMOOD N., BLACK M. J.: Breathing life into shape: Capturing, modeling and animating 3D human breathing. *ACM Transactions on Graphics* 33, 4 (2014), 52:1–52:11.
- [Tri14] Trimble: Trimble 3D warehouse. <http://sketchup.google.com/3Dwarehouse/>.
- [TV08] TANGELDER J. W., VELTKAMP R. C.: A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications* 39, 3 (2008), 441–471.
- [TYK\*12] TALTON J., YANG L., KUMAR R., LIM M., GOODMAN N., MĚCH R.: Learning design patterns with Bayesian grammar induction. In *Proceedings of UIST* (Cambridge, MA, USA, 2012), pp. 63–74.
- [Val84] VALIANT L. G.: A theory of the learnable. *Communications of the ACM* 27, 11 (1984), 1134–1142.
- [vKTS\*11] VAN KAICK O., TAGLIASACCHI A., SIDI O., ZHANG H., COHEN-OR D., WOLF L., HAMARNEH G.: Prior knowledge for part correspondence. *Computer Graphics Forum* 30, 2 (2011), 553–562.
- [vKXZ\*13] VAN KAICK O., XU K., ZHANG H., WANG Y., SUN S., SHAMIR A., COHEN-OR D.: Co-hierarchical analysis of shape structures. *ACM Transactions on Graphics* 32, 4 (2013), 69:1–10.

- [VKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707.
- [WAVK\*12] WANG Y., ASAFI S., VAN KAICK O., ZHANG H., COHEN-OR D., CHEN B.: Active co-analysis of a set of shapes. *ACM Transactions on Graphics* 31, 6 (2012), 165:1–165:10.
- [WBLP11] WEISE T., BOUAZIZ S., LI H., PAULY M.: Realtime performance-based facial animation. *ACM Transactions on Graphics* 30, 4 (2011), 77:1–77:10.
- [WGW\*13] WANG Y., GONG M., WANG T., COHEN-OR D., ZHANG H., CHEN B.: Projective analysis for 3D shape segmentation. *ACM Transactions on Graphics* 32, 6 (2013), 192:1–192:12.
- [WHG13] WANG F., HUANG Q., GUIBAS L.: Image co-segmentation via consistent functional maps. In *Proceedings of ICCV* (Sydney, Australia, 2013).
- [WLVGP09] WEISE T., LI H., VAN GOOL L., PAULY M.: Face/off: Live facial puppetry. In *Proceedings of Symposium on Computer Animation* (Copenhagen, Denmark, 2009), pp. 7–16.
- [WS13] WANG L., SINGER A.: Exact and stable recovery of rotations for robust synchronization. *Information and Inference* 2, 2 (2013), 145–193.
- [WSK\*15] WU Z., SONG S., KHOSLA A., YU F., ZHANG L., TANG X., XIAO J.: 3D shapenets: A deep representation for volumetric shapes. In *Proceedings of CVPR* (Boston, MA, USA, 2015).
- [XCF\*13] XU K., CHEN K., FU H., SUN W.-L., HU S.-M.: Sketch2scene: Sketch-based co-retrieval and co-placement of 3D models. *ACM Transactions on Graphics* 32, 4 (2013), 123:1–12.
- [XHS\*15] XU K., HUANG H., SHI Y., LI H., LONG P., CAICHEN J., SUN W., CHEN B.: Autoscanning for coupled scene reconstruction and proactive object analysis. *ACM Transactions on Graphics* 34, 6 (2015), 177:1–177:14.
- [XKHK14] XU K., KIM V., HUANG Q., KALOGERAKIS E.: Wikipage: Data-driven shape analysis and processing. <http://wp.cs.umass.edu/datadrivenshape/>.
- [XLZ\*10] XU K., LI H., ZHANG H., COHEN-OR D., XIONG Y., CHENG Z.-Q.: Style-content separation by anisotropic part scales. *ACM Transactions on Graphics* 29, 5 (2010), 184:1–184:10.
- [XMZ\*14] XU K., MA R., ZHANG H., ZHU C., SHAMIR A., COHEN-OR D., HUANG H.: Organizing heterogeneous scene collections through contextual focal points. *ACM Transactions on Graph.* 33, 4 (2014), 35:1–12.
- [XSX\*14] XU W., SHI Z., XU M., ZHOU K., WANG J., ZHOU B., WANG J., YUAN Z.: Transductive 3D shape segmentation using sparse reconstruction. *Computer Graphics Forum* 33, 5 (2014), 107–115.
- [XXLX14] XIE Z., XU K., LIU L., XIONG Y.: 3D shape segmentation and labeling via extreme learning machine. *Computer Graphics Forum* 33, 5 (2014), 85–95.
- [XXM\*13] XIE X., XU K., MITRA N. J., COHEN-OR D., GONG W., SU Q., CHEN B.: Sketch-to-design: Context-based part assembly. *Computer Graphics Forum* 32, 8 (2013), 233–245.
- [XXS\*15] XIE Z., XU K., SHAN W., LIU L., XIONG Y., HUANG H.: Projective feature learning for 3D shapes with multi-view depth images. *Computer Graphics Forum* 34, 7 (2015), 1–11.
- [XZCOC12] XU K., ZHANG H., COHEN-OR D., CHEN B.: Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Transactions on Graphics* 31, 4 (2012), 57:1–57:10.
- [XZZ\*11] XU K., ZHENG H., ZHANG H., COHEN-OR D., LIU L., XIONG Y.: Photo-inspired model-driven 3D object modeling. *ACM Transactions on Graphics* 30, 4 (2011), 80:1–80:10.
- [YCHK15] YUMER M. E., CHAUDHURI S., HODGINS J. K., KARA L. B.: Semantic shape editing using deformation handles. *ACM Transactions on Graphics* 34 (2015), 86:1–86:12.
- [YCM14] YUMER M. E., CHUN W., MAKADIA A.: Co-segmentation of textured 3D shapes with sparse annotations. In *Proceedings of CVPR* (Columbus, OH, USA, 2014).
- [YK12] YUMER M. E., KARA L. B.: Co-abstraction of shape collections. *ACM Transactions on Graphics* 31, 6 (2012), 166:1–166:11.
- [YK14] YUMER M., KARA L.: Co-constrained handles for deformation in shape collections. *ACM Transactions on Graphics* 32, 6 (2014), 187:1–187:11.
- [YN10] YU K., NG A.: Feature learning for image classification. In *ECCV Tutorials* (2010).
- [ZF14] ZEILER M. D., FERGUS R.: Visualizing and understanding convolutional neural networks. In *Proceedings of ECCV* (2014), 187:1–187:1.
- [ZKP10] ZACH C., KLOPSCHITZ M., POLLEFEYS M.: Disambiguating visual relations using loop constraints. In *Proceedings of CVPR* (Colorado Springs, CO, USA, 2010).
- [ZMT05] ZHANG E., MISCHAIKOW K., TURK G.: Feature-based surface parameterization and texture mapping. *ACM Transactions on Graphics* 24, 1 (2005), 1–27.
- [ZSSS13] ZIA M. Z., STARK M., SCHIELE B., SCHINDLER K.: Detailed 3D representations for object recognition and modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 11 (2013), 2608–2623.
- [ZWK14] ZHAO X., WANG H., KOMURA T.: Indexing 3D scenes using the interaction bisector surface. *ACM Transactions on Graphics* 33, 3 (2014), 22:1–22:15.
- [ZXTZ14] ZHANG Y., XU W., TONG Y., ZHOU K.: Online structure analysis for real-time indoor scene reconstruction. *ACM Transactions on Graphics* 34, 5 (2014), 159:1–159:12.