Int J Comput Vis
DOI 10.1007/s11263-009-0279-0

Contextual Part Analogies in 3D Objects

L. Shapira · S. Shalom · A. Shamir · D. Cohen-Or · H. Zhang

Received: 17 October 2008 / Accepted: 21 July 2009

Abstract In this paper we address the problem of finding analogies between parts of 3D objects. By partitioning an object into meaningful parts and finding analogous parts in other objects, not necessarily of the same type, many analysis and modeling tasks could be enhanced. For instance, partial match queries can be formulated, annotation of parts in objects can be utilized, and modeling-by-parts applications could be supported. We define a similarity measure between two parts based not only on their local signatures and geometry, but also on their context within the shape to which they belong.

In our approach, all objects are hierarchically segmented (e.g. using the shape diameter function), and each part is given a local signature. However, to find corresponding parts in other objects we use a context enhanced part-in-whole matching. Our matching function is based on bipartite graph matching and is computed using a flow algorithm which takes into account both local geometrical features and the partitioning hierarchy. We present results on finding part analogies among numerous objects from shape repositories, and demonstrate sub-part queries using an implementation of a simple search and retrieval application. We also demonstrate a simple annotation tool that carries textual tags of object parts from one model to many others using analogies, laying a basis for semantic text based search.

Keywords Part retrieval · Shape signature · Hierarchical partitioning · Distance measure · Shape matching

L. Shapira (✉) · D. Cohen-Or
Tel-Aviv University, P.O. Box 39040, Tel-Aviv 69978, Israel
e-mail: liors@post.tau.ac.il

D. Cohen-Or
e-mail: cohenor@gmail.com

S. Shalom
Bar-Ilan University, Ramat Gan 52900, Israel
e-mail: shooshx@gmail.com

A. Shamir
Interdisciplinary Center, Herzliya, Israel
e-mail: arik.shamir@gmail.com

H. Zhang
Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada
e-mail: haoz@cs.sfu.ca

1 Introduction

Adding semantic information to 3D models is beneficial for many analysis and modeling tasks. For instance, recognizing a specific functional part in a model provides the ability to search for analogous parts for modeling applications, comparative studies, etc. However, unlike documents retrieval, it is still almost impossible to search parts of 3D objects semantically or using text tags. The main reason for this is that most digital 3D models are not partitioned into semantic parts nor annotated.

Usually, the search for 3D objects is based on finding objects similar to a given query model with the aid of some global shape signature. Searching for a specific part inside a 3D model is more challenging when in most cases the models are not segmented and the different parts composing an object are not linked to semantic tags. We present a method that finds analogies among parts of digital 3D models by segmenting them and creating a *contextual* signature for each part. That is, a part is not only characterized by its own geometric properties but also its context within the whole shape.

Published online: 05 August 2009

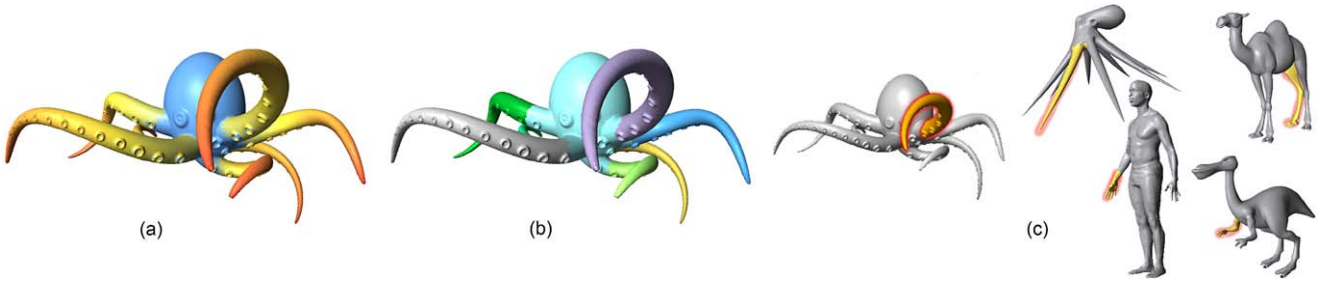

Fig. 1 (Color online) Contextual part analogies: (**a**) The SDF is calculated on all objects in the database (colored from *red* for narrow diameter, to *blue* for wide diameter). (**b**) All objects are then partitioned hierarchically and the signatures of their parts is stored. (**c**) Our similarity measure finds analogies between similar parts of different objects. For example, analogous parts to the octopus arm are found in different models even if the parts and the whole objects are both dissimilar in their shapes. Signatures that are based on the geometry of the parts alone can not find such analogies. When we factor in the context of the part, these similarities emerge

The dominating representation of digital 3D objects is a 2D boundary surface mesh. However, when studying a segmentation of models into parts, and analogies between different parts, one notes that many of the cues to a good partitioning, and natural analogies, are *volumetric* in nature. Therefore, it is beneficial to base the segmentation and analogies on volumetric attributes. In this paper we use the shape-diameter function (SDF, Shapira et al. 2008) for segmentation and also as the basis for two of the local shape signature measures we use. The SDF provides a link between the object's volume and the mesh surface, mapping volumetric information onto the surface boundary mesh. It is defined by examining the diameter of the model in the neighborhood of each point on its boundary surface. To define a *contextual* signature for each part, there is a need to define the relation between the part and the whole shape. Hence, we extend (Shapira et al. 2008) partitioning algorithm to create a hierarchical segmentation of the objects storing sub-parts contained in a given part as its child nodes in a tree representation of the object.

To find analogies in a database of objects, we first partition all given objects into parts. Next, we define a signature for each part based on geometric attributes as well as its relation to the whole object (its context). The context is defined by the path from the part to the root of the object's partitioning. Given two parts we define a context-aware distance measure between them by using bi-partite graph matching between the two characteristic paths of the parts.

When a user specifies a model part query, we can retrieve the most similar parts from all models in the database based on the context-aware similarity measure. Signatures that are based only on the geometry will fail to find many similarities (Fig. 1). This can be utilized, for instance, to carry tags from an annotated part to all similar parts in a database. Later, these tags can be used for text-based retrieval from the database.

Our main contributions in this paper are therefore:

1. An automatic, consistent and hierarchical clustering algorithm based on the SDF, improving the algorithm of Shapira et al. (2008).
2. A novel similarity measure which takes into account not only local shape descriptions, but is also context-aware and provides the ability to find part analogies.
3. We demonstrate that our measure can work with many local distance measure, and with any database containing hierarchically segmented models.
4. We demonstrate applications of part analogies in part-in-whole queries and partial matching, as well as for a tagging and annotation tool for parts in a 3D shape database.

2 Related Work

There are numerous mesh partitioning techniques based on various mesh attributes. For a survey on different mesh partitioning techniques we refer the reader to (Shamir 2007). Since we seek part-type partitioning, we employ the method in (Shapira et al. 2008), which uses the SDF to partition sets of objects consistently. This measure relates to the medial axis transform (MAT) (Choi et al. 1997) which is extremely informative for shape analysis and partitioning. The SDF replaces the local shape radius of the MAT by a measure of the local *shape diameter*. Partitioning based on the SDF is likely to create parts which are similar in their SDF signature and are correspondent among different objects (Fig. 2). Nevertheless, other part-type partitioning methods such as Attene et al. (2006b), Katz and Tal (2003) could also be used in the first stage of our algorithm.

Shape matching is also an active topic of research, and numerous shape signatures based on geometry and topology have been proposed (Bustos et al. 2005; Gal et al. 2007; Tangelder and Veltkamp 2008). In Osada et al. (2002), a large number of points are sampled from the surface and several different statistics are gathered. These statistics form the basis for several histogram based signatures which are

Fig. 2 Sample models used in this paper. (**a**) Models from the SHREC model set (Giorgi et al. 2007), which contains around 400 categorized models. (**b**) Models from the PSB (Shilane et al. 2004) and other various sources (over 300 models). The color mapping on the models is of the Normalized SDF values. Note how these values already indicate similarity between analog parts

simple, yet descriptive. In Gal et al. (2007) the SDF was used alongside a centricity measure to define a 2D histogram based signature which is pose-oblivious. In Ben-Chen and Gostman (2008) a signature based on conformal geometry is proposed, it is invariant under non-rigid quasi-isometric transformations. In Sect. 4 we evaluate these measures, in addition to a simple SDF histogram based signature, and use them as basis for a novel, contextual similarity measure.

We concentrate on matching sub-parts of objects which is related more to partial matching. Nevertheless, partial matching techniques (Belongie et al. 2001; Funkhouser and Shilane 2006; Gal and Cohen-Or 2006; Johnson and Hebert 1999; Novotni et al. 2005; Zhang et al. 2008) are based more on feature correspondence but less on segmentation results. Segmentation and shape analysis was used to enrich models with semantic information in Attene et al. (2006a, 2007) by manually connecting parts through a user interface to an instance in a knowledge base, or by using ontology connecting form and functionality in Camossi et al. (2007).

Finding correspondence and analogies between different shapes is recently becoming an active field of research. Many works focus on the need to define measures for similarity of shapes (Sundar et al. 2003; Cornea et al. 2005; Funkhouser et al. 2005), alignment of shapes (Gelfand et al. 2005) and finding complete correspondence between two models for the purposes of deformations, morphing (Alexa 2002), and cross-parameterizations (Kraevoy and Sheffer 2004). We focus on finding analogies between sub-parts.

Low level analogies to match vertices are used in Schreiner et al. (2004), Kraevoy and Sheffer (2004) for cross parametrization between two models. The matching is found using user supplied matching points. In Sumner and Popovic (2004) such matching points are used to define a many-to-

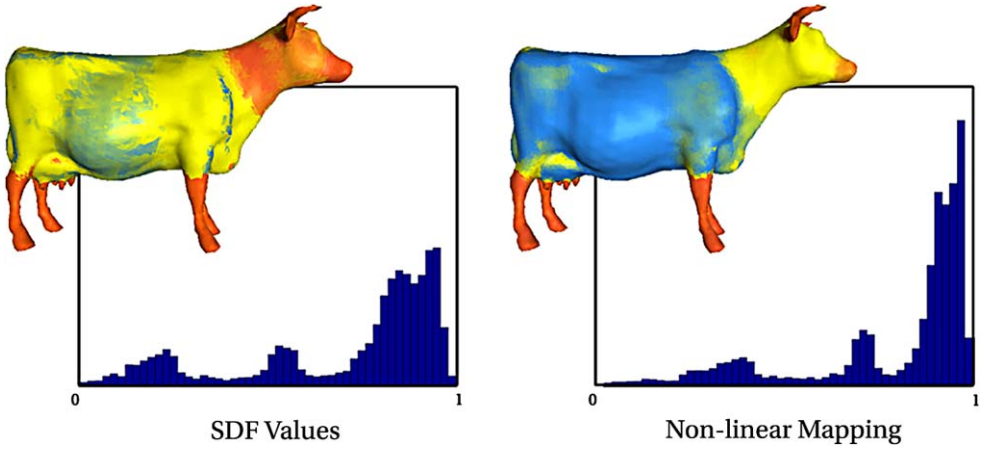

Fig. 3 Non-linear (log-space) mapping of SDF values to enhance the importance of delicate parts of the model (low SDF values). Note how in log-space the horns and nose of the cow are better separated from the head

many mapping of vertices between the models which is then used to transfer deformations between the models. These works require a long time to process and non-trivial user input. They operate mostly on two models, and cannot work on parts of models. Our work targets higher level analogies and can work also on objects with different topology and structure (Fig. 1).

In our experiments we show results on models from the Princeton Shape Benchmark (PSB) (Shilane et al. 2004) and the SHREC water-tight models database (Giorgi et al. 2007) (Figs. 2, 9).

3 Partitioning to Parts

When examining 3D models, one can observe that the similarity of parts often stems from their functionality. For example in humans and animals parts are associated with organs, which are usually 3D *volumetric* sub-parts of the shape. Therefore, an automatic algorithm aimed at detecting such 3D shape analogies must first identify these sub-parts. We use a partitioning of the shapes guided by the *shape diameter function* (SDF) (Shapira et al. 2008). The SDF connects volumetric information of an object onto its boundary mesh by measuring the local diameter of the object at points on its boundary. Hence, the SDF is suitable to guide volumetric part extraction, detect natural 3D shape partitioning, and define part signatures (Fig. 2).

The SDF at a point on the surface of the object is defined as the diameter of the object in the neighborhood of that point. Given a point on the surface mesh a set of rays is sent inside a cone centered around its inward-normal direction (the opposite direction of its normal) to the other side of the mesh. Ideally we would use only one ray, opposite the normal. However, in order to obtain a smooth function and better approximate the shape diameter in the presence of geometric noise, we must sample several rays. The value of the SDF at the point is defined as a weighted average of all the lengths of the rays that fall within one standard deviation from the median of all lengths. We use the inverse of the angles between the rays to the center axis of the cone as the weight to put more emphasis on rays opposite the normal direction.

To maintain compatibility over different meshes, which may have different scales and resolutions, we normalize and smooth the SDF values. We also perform the partitioning in log-space to enhance the importance of delicate parts, which tend to have low characteristic SDF values (Fig. 3). The normalized SDF value $nsdf$ of facet f is calculated at the centroid of the face and defined as

$$nsdf(f) = \log\left(\alpha \cdot \frac{sdf(f) - \min(sdf)}{\max(sdf) - \min(sdf)} + 1\right) \Big/ \log(\alpha + 1),$$

where $sdf : F \to \mathbb{R}$ is the SDF value for each facet f and α is a normalizing parameter which is set to 4 in all our examples.

The SDF can be seen as a scalar function over the mesh surface. Specific iso-values of the SDF create iso-contours on the surface, which can be used to partition the mesh. The partitioning algorithm consists of two steps, first we model the SDF values and then we cluster the faces of the mesh.

Partitioning Algorithm First, we use a Gaussian Mixture Model (GMM) to fit k Gaussians to the histogram of SDF values of the faces. This is achieved using the Expectation-Maximization (EM) (Bilmes 1997) algorithm (Fig. 4). Once we have a GMM, we calculate for each face f, a vector $v_f \in R^k$ where v_f^i is the probability of f to belong to the ith Gaussian. Therefore at this stage, each face on the mesh, belongs to all Gaussians with some probability. In the following step we use this information to cluster mesh faces together, and create a segmentation.

In the second step we would like to segment the model into parts. Additionally, we would like to ensure that the boundaries between parts adhere to local mesh features such as concave areas or creases and are smooth. We employ an

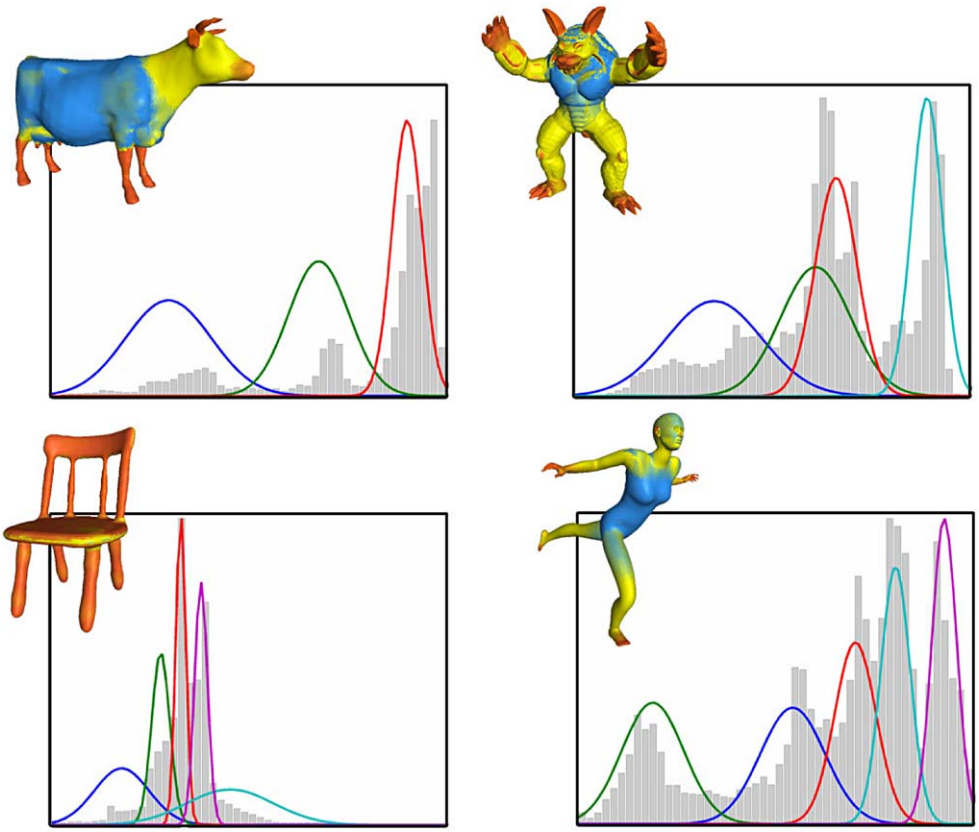

Fig. 4 GMM model on normalized SDF values calculated using EM. The GMM is later used to partition the model. For illustrative purposes we show here a varying number of Gaussians on the different models. Note that when applying automatic partitioning to a whole database of models, we used a constant number of Gaussians (4), with a handful of exceptions

alpha expansion graph-cut algorithm (Boykov et al. 2001; Zabih and Kolmogorov 2004) to solve the k-way graph-cut problem, leading to a labeling of the mesh faces. The graph-cut problem is known to be NP-hard. The alpha expansion algorithm utilizes a series of large moves, changing a large number of mesh face labels at a time, to arrive at an approximate solution within a known factor of the optimal solution.

We define a set of k labels, such that label i corresponds to cluster i from the GMM. Let us denote by $\hat{x} : F \to B$, the face labeling, where F is the set of mesh faces and B is the label set. When optimizing for $\hat{x}$, we wish to take into account both the cluster assignment probabilities computed from the EM step and the quality of the boundaries. Therefore, our graph-cut formulation minimizes the following energy functional composed of two terms: e_1, a *data term*, and e_2, a *smoothness term* (see also Katz and Tal 2003; Shapira et al. 2008).

$$\mathcal{E}(\hat{x}) = \sum_{f \in F} e_1(f, \hat{x}(f)) + \lambda \sum_{\{f,g\} \in N} e_2(\hat{x}, f, g),$$

$$e_1(f, b) = -\log(P(f|b) + \varepsilon),$$

$$e_2(\hat{x}, f, g) = \begin{cases} l(f,g)(1 - \log(\theta(f,g)/\pi)), & \hat{x}(f) \neq \hat{x}(g), \\ 0, & \hat{x}(f) = \hat{x}(g) \end{cases}$$

where

- $P(f|b)$ represents the probability of assigning face f to cluster b; these values are derived from the EM-fitted GMM in the first step of the algorithm;
- $\theta(f, g)$ is the dihedral angle between facets f and g (if connected, see next paragraph);
- $l(f, g)$ is the length of the edge shared by f and g;
- N is the set of adjacent face pairs in the mesh.

A constant value of $\lambda = 0.3$, as a weight for smoothing, gave good results in all our experiments. We also normalize smoothness by the edge length $l(f, g)$. Here and in subsequent equations, $\varepsilon = 10^{-3}$ is used to avoid numerical instability.

A large percentage of the models used in the paper are challenging: many of them are not water-tight, contain inner parts, have faulty connectivity etc. We overcome these problems by utilizing a spatial search structure to find neighboring faces (i.e. not entirely relying on connectivity). This allows us to support a wide variety of models. Note that in the smoothness term we use the dihedral angle between two connected faces. If two faces are adjacent but are not connected, we use instead a constant value, which we have found to produce good results.

The result of the graph-cut algorithm is a labeling of the mesh faces, where each label corresponds to a Gaussian in the model (see Fig. 5). These label are the basis of the hierarchical segmentation of the model. In the next subsec-

tion we describe how the association of each face with the k Gaussians is used to create a this hierarchical partitioning. Using more Gaussians in the mixture creates a finer segmentation of the mesh into parts, and also increases the number of levels in the hierarchy of sub-parts. We have found that for most models we have used, a constant value of $k = 4$ gave good results. A few exceptions were re-partitioned using a different value (between 3 and 6). Figure 6 demonstrate that our scheme is not highly dependant on the choice of parameter values. We show partitioning results for the Armadillo varying both k and λ parameter values. Note how the partitioning does not change drastically, remaining consistent throughout the different values.

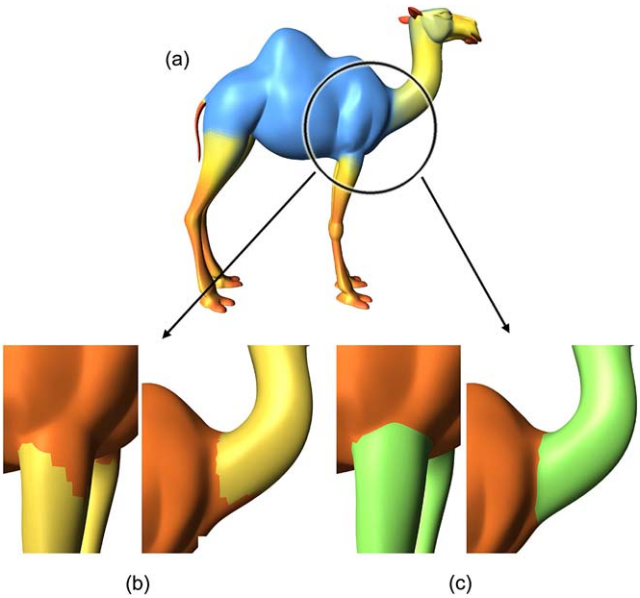

Fig. 5 Natural part boundaries. (**a**) Camel model with SDF visualization. (**b**) Selecting the label matching the closest Gaussian for each face partitions the mesh without adhering to local mesh features. (**c**) Applying the graph-cut step smooths the boundaries

3.1 Hierarchical Partitioning

Many times analogies between parts are based on the relation of the parts to their respective whole objects. For example the leg on a human model, if seen out of context, resembles a cylinder. However, seen in context, it is analogous to the legs of other bipeds and quadrupeds. Parts from different objects that vary in their geometric shape or attributes individually, become analogous when placed in the context of their whole shape. Therefore, we want to create a hierarchical representation of each shape's parts, and employ it later to find analogies.

We sort the means of the GMM model from large to small, and define k' iso-values separating the Gaussians, and consequently, separating the mesh into "levels". The first level is always set at 1 and is considered the root of the object's partitioning hierarchy, represented by a tree. The next value separates the label corresponding to the Gaussian with the largest mean from the rest of the labels. Thus, each face of the mesh is assigned one of two possible labels. For example, in a human model this would separate the torso from the head and limbs. The actual parts are defined by propagating from a seed triangle to create connected pieces of the model. The next level separates the label correspond-

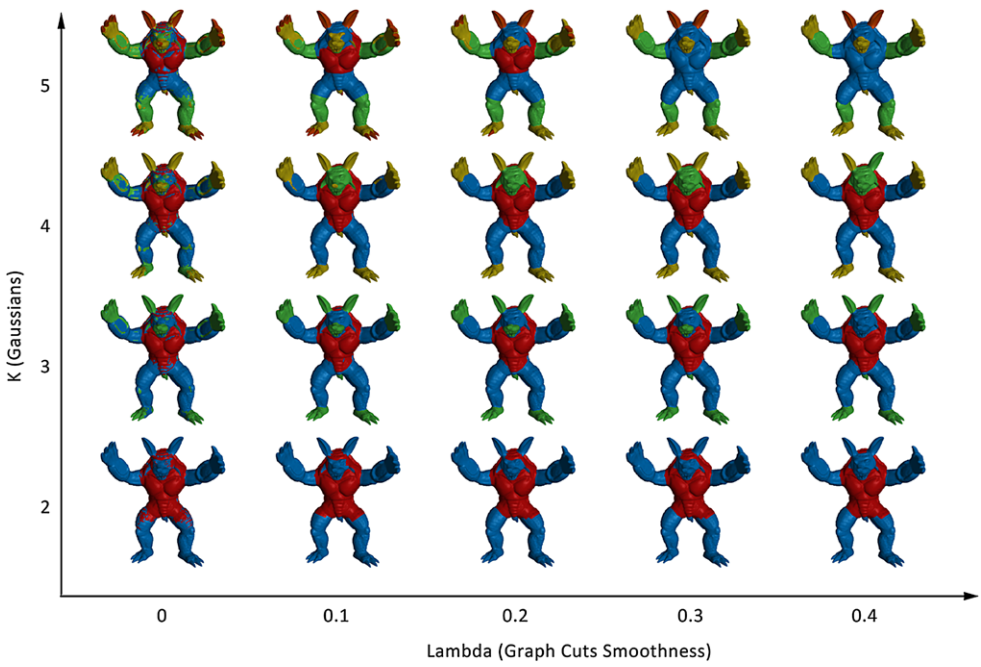

Fig. 6 Our automatic segmentation of the Armadillo model over a wide range of parameters creates proper segmentation which is also consistent

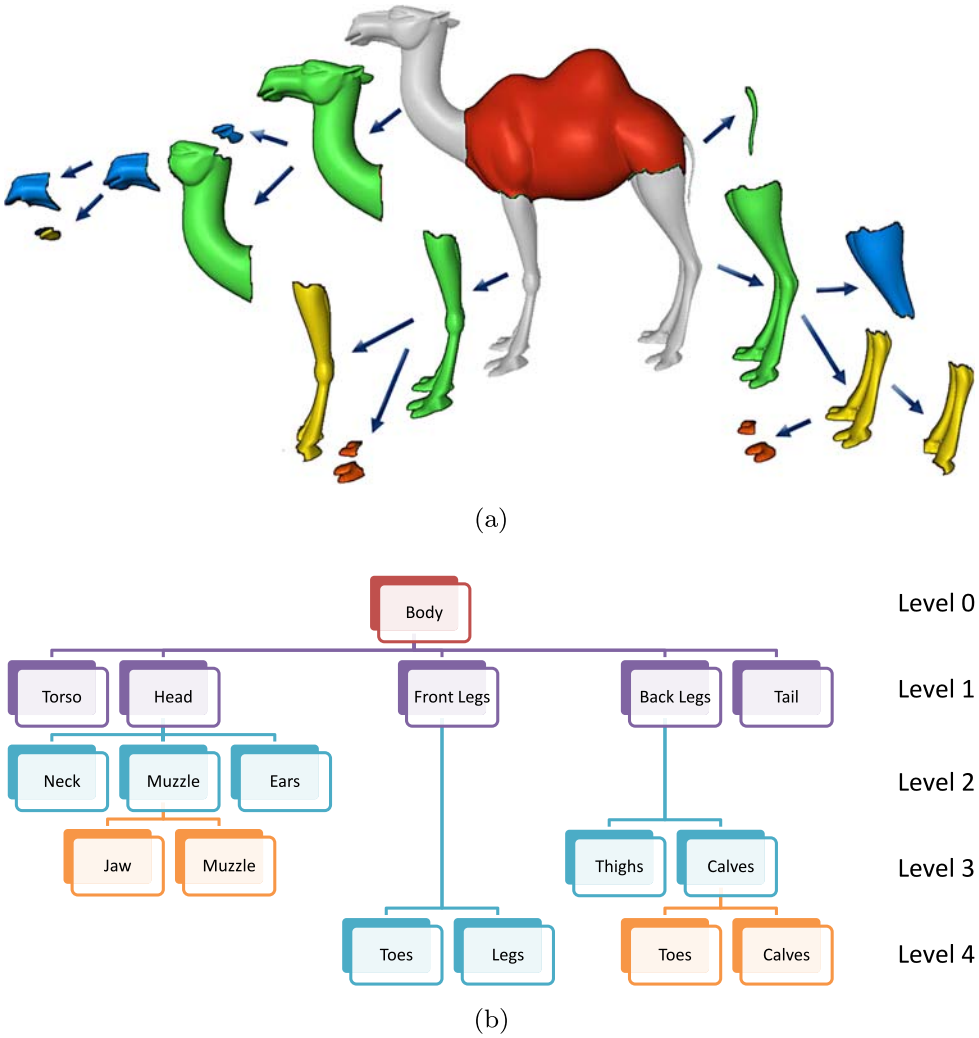

Fig. 7 Hierarchical partitioning. (**a**) The camel model is partitioned automatically using four iso-values, resulting in a five-tier hierarchy of partitioning. (**b**) The partitioning induces a hierarchical graph of parts

ing the Gaussian with the second largest mean, separating the faces of the object into three distinct groups. For example, in a human model this would now separate the hands from the arms and the feet from the legs. The rest of the iso-values are calculated similarly in a recursive manner. Iso-values are merged when they do not change the partitioning, thus $k' \leq k$. For example, if we use $k = 2$ then each part is associated with one of two Gaussian means. The first hierarchy level will be at iso-value 1 (the whole model), and the second iso-value would be between the Gaussian means.

Therefore, given a set of iso-values, for each iso-value, each face in the model belongs to a certain part. Using this hierarchy we build the segmentation tree of the model. For example, the five-tier camel partitioning hierarchy (Fig. 7(a)) induces a hierarchical part graph as can be seen in Fig. 7(b). Note that the toes and legs of the camel (in all four legs) were separated using the fourth iso-value. However, in the front, they are direct descendants of the front legs, while in the back, the third iso-value induces a slightly more detailed partitioning. In any case, once we build the partitioning hierarchy there is no need to remember the specific iso-values. Additional examples of hierarchical segmentations can be seen in Fig. 8. The tree defines the relation of parts inside the object, and assists to define a better distance metric to recognize similar object parts as described in the next section.

4 Part Signature and Distance Measure

To find analogies between multiple models, we must define a way to measure similarity between parts of models. We contend that when seeking to compare two parts, the context from which they came is crucial to the comparison. A finger on a human model is just a capped cylinder. However, when taken in context of the hand, the arm and the entire body, its description is more complete, and better matches and analogies could be found.

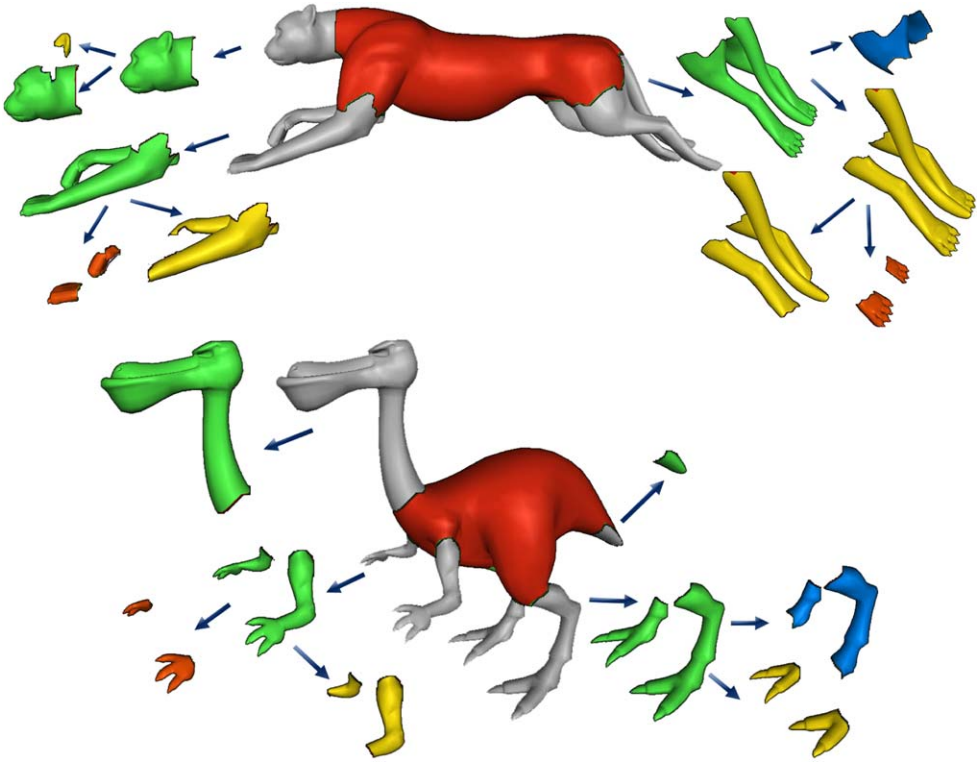

Fig. 8 Hierarchical partitioning of the cheetah and dinopet models

Each segmented part in the model is assigned a local signature composed of its geometrical attributes. For the purpose of this work we have experimented with the following signatures (see also Fig. 10):

- *HSDF*: Normalized histogram of SDF values within the part and the size of the part as a percentage of the whole model (see Sect. 4.1).
- *SD (D1, D2, D3, A3)*: Shape-distribution signatures from (Osada et al. 2002). A large number of points are sampled uniformly on the surface of the model. The signature is a histogram on the values of the following functions: (D1) Distances between a fixed point and the sampled points. (D2) Distances between two random points on the surface. (D3) The square root of the area of the triangle between three random points. (A3) The angle between three random points.
- *CG*: Conformal geometry signatures from (Ben-Chen and Gostman 2008), which define a curvature based histogram measure.

In some cases, the geometrical attributes are sufficient to define a good distance measure between the parts, specifically in distinct parts, such as a head or paw (Fig. 11). However, analogies may stem from the characteristic of the part in the whole as well as its geometric attributes. Therefore, we define a context-based similarity measure, which uses both local distance measures, and part-in-whole information gathered from the hierarchical partitioning of the model. We show, through experimentation, that this measure improves the results of all the local distance measures we tried. Moreover, the context based is robust to different partitioning results, enhancing its usability.

Using the hierarchy we define the context of a part as the path between the node representing the part, and the root of the partitioning hierarchy. Each node along this path represents a part for which we can calculate the geometrical attributes as described above. The set of all of these geometric attributes define the context descriptor of the part. We use this context descriptor in a distance measure between two parts, that takes into account both the similarity of the parts themselves, and the similarity among the path nodes.

4.1 Local SDF Signature

We define the local *HSDF* measure between two parts p and p' as a weighted sum of the distance between the local part histograms, and the relative part sizes.

$$d_{histogram}(p, p') = \left\| H(p)/\|H(p)\| - H(p')/\|H(p')\| \right\|_2$$

$$d_{size}(p, p') = \frac{|size(p) - size(p')|}{size(p) + size(p')},$$

$$HSDF(p, p') = \frac{1}{3} \cdot d_{size}(p, p') + \frac{2}{3} \cdot d_{histogram}(p, p').$$

H is a normalized histogram of the part's SDF values. We construct the histogram based on the original SDF mea-

Fig. 9 Partitioning to parts using the SDF results works on a wide variety of models, as seen here for sample models from the SHREC model set (**a**), the PSB, and other sources (**b**). All the results seen above were achieved automatically with a small (4) number of Gaussians fitted to SDF histograms

surements in the part, removing the top and bottom 5% to remove possible outliers. We use the *L1* distance metric on the normalized histogram, treating it as a vector. We have experimented with various distance measures such as *L2*, *Chi-Square* and *Kullback-Leibler* (1), but found that results do not differ significantly from *L1*. In the above expressions $size(p)$ is the relative size of part p within the whole shape.

4.2 Context-Aware Distance Measure

We define a local geometric distance measure between parts p and p' as $d(p, p')$. In our experiments we have used HSDF, CG, SD(D1), SD(D2), SD(D3), SD(A3) as local part signatures.

To define our context-aware distance measure $D(p, p')$ we consider not only the local distance measure $d(p, p')$ between the two parts, but also the whole *paths* from the nodes of p and p' to the root of their partitioning hierarchy. Given two such paths on which we want to measure similarity, we build a bipartite graph G (Fig. 12) such that each side represents all nodes in each path. The edges between the two sides contain one edge between p and p' (the two parts whose distance we want to measure), and an edge between each ancestor of p to each ancestor of p'. Note that the number of ancestors of p and p' may be different. The capacity of an edge between two nodes q and q' is defined as:

$$capacity(q, q') = \frac{1}{d(q, q') + \varepsilon} - 1$$

Lastly, we add two nodes, source S and sink T, and connect each one of them to the nodes in one side of the graph re-

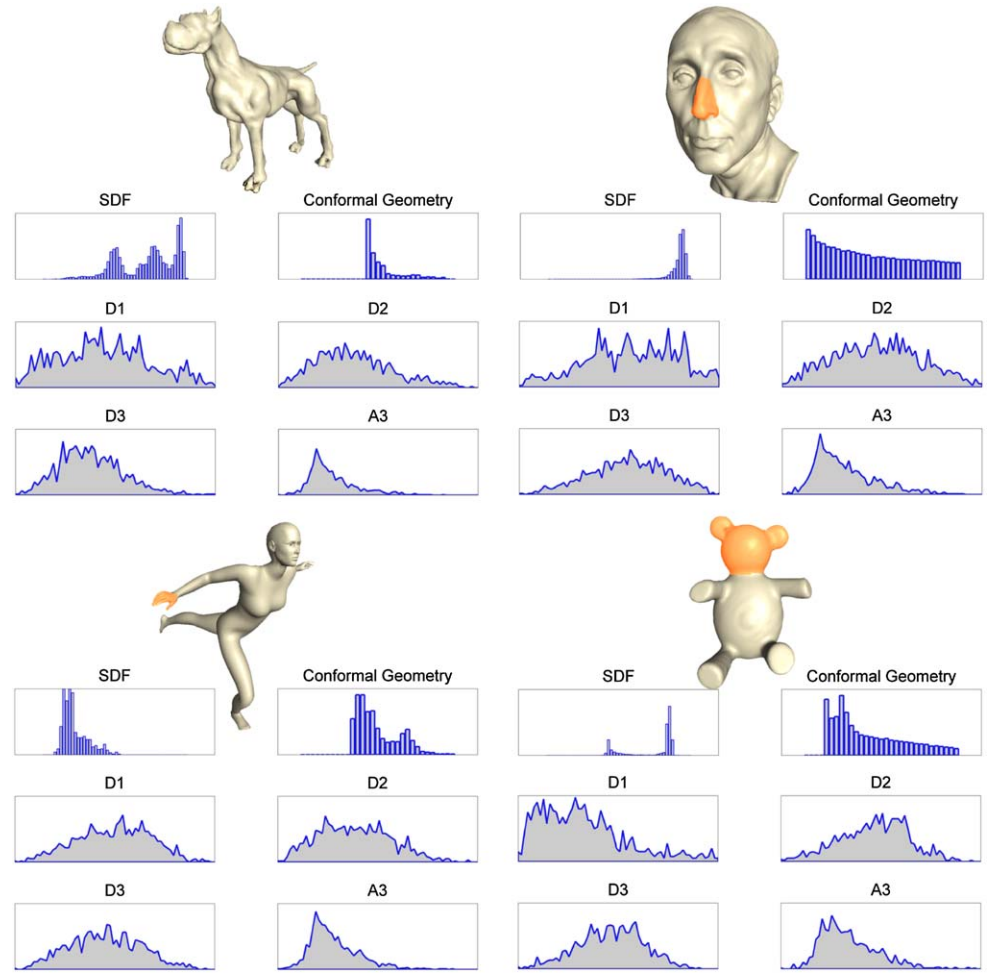

Fig. 10 Different local shape signatures, as visualized here for the dog (whole model), nose of statue, hand of woman and head of teddy bear. SDF is a histogram of normalized SDF values. Conformal Geometry (CG) is a histogram of sampled values as described in Ben-Chen and Gostman (2008). D1,D2,D3 and A3 are histograms of different measures as described in Osada et al. (2002)

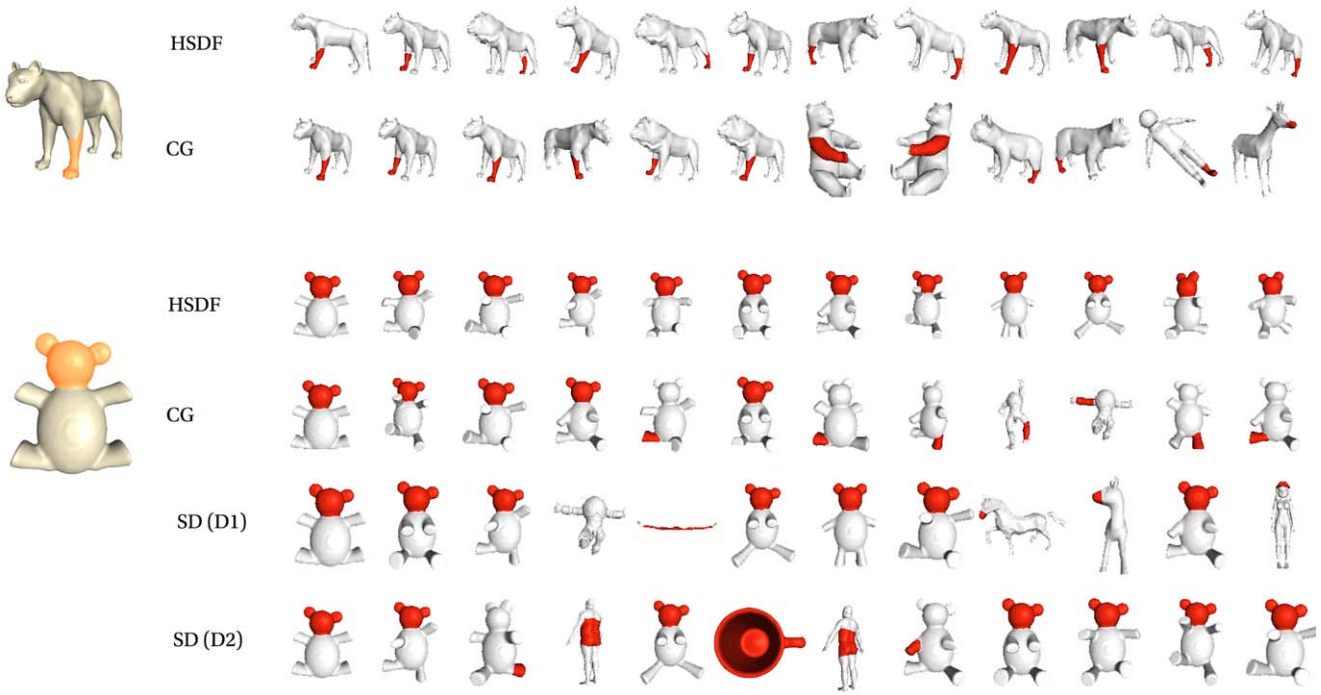

Fig. 11 Local geometrical attributes in some cases are sufficient to define a good distance measure, as evident in the cougar paw and the teddy bear's head. Note that success rates vary with the choice of a specific distance measure

spectively, with capacity equal to $\beta \cdot capacity(p, p')$, with β set at 1.5. This serves as an upper limit on the capacity of the flow in the graph G. We now define

$$D(p, p') = \frac{1}{flow(G) + 1}$$

where $flow(G)$ is the maximum flow in graph G.

The key motivation behind such a measure is on one hand to match the part in the context of the whole hierarchy, and on the other to achieve robustness against differences in the partitioning. The measure will be higher as more parts in the path from the node to its root match the respective nodes in the compared part. However, it is hard to determine the exact matching of parts in two hierarchies. By connecting each ancestor of p to each ancestor of p' we are assured that the flow will represent the maximum similarity from possible different matchings. For instance, given three geometrically similar parts p, q, w such that p and q also come from similar models, the local distance $d(p, q)$ and $d(p, w)$ will be similar. When constructing the context-aware distance measure, we construct two graphs, one comparing p to q, the other comparing p to w. In these graphs, the edges pq and pw (defined using $d(p, q)$ and $d(p, w)$ respectively) have similar weights. However, when connecting nodes along the path from p to its hierarchy root to the path of q and the path of w, we get different edge weights (since p and q originate from similar models, with similar parts). Therefore $D(p, q) < D(p, w)$. Figure 13 illustrates this idea.

Comparing a local signature to a context-aware distance measure (using that same local signature for comparing specific parts) shows significant advantage to using the context-aware measure. Such examples can be seen in Fig. 14.

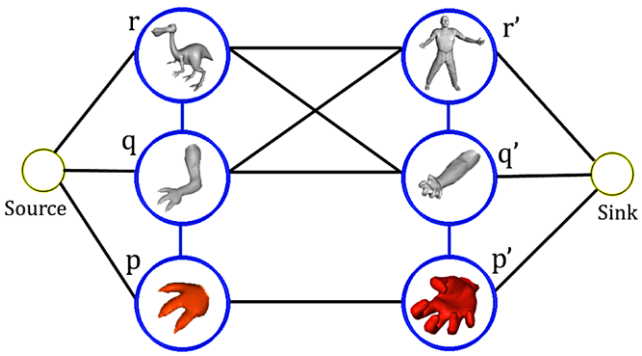

Fig. 12 Measuring context based similarity between two parts using a bipartite graph. The first part hierarchy is represented by the nodes p, q, r while the second part hierarchy is represented by the nodes p', q', r'. The capacity of the edge (x, y) is defined to be $\frac{1}{d(x,y)+\varepsilon}$. The similarity between two parts is defined as the maximum flow through this graph

5 Applications and Results

We will demonstrate the usefulness of our part analogies approach using two applications. The first is in the context of a search and retrieval application of 3D shapes. Using analogies, one can search for parts of shapes in the database that are similar to a given part, or for objects that contain similar parts to a given part query. The second can be seen as a tool to enhance the meta-data in 3D objects. Once part analogies are found, any information linked with the query part can be carried automatically to other anal-

Table 1 Example of local vs. context-aware distance measures between the dinopet's hand and six other parts (Fig. 13), where the *HSDF* distance measure is used

Part	$d(p, p')$	$D(p, p')$
a (dinopet other hand)	0.001	0.0033
b (human hand)	0.034	0.0147
c (dinosaur hand)	0.042	0.0263
d (cheetah paw)	0.07	0.0242
e (human head)	0.26	0.126
f (airplane wing)	0.373	0.192

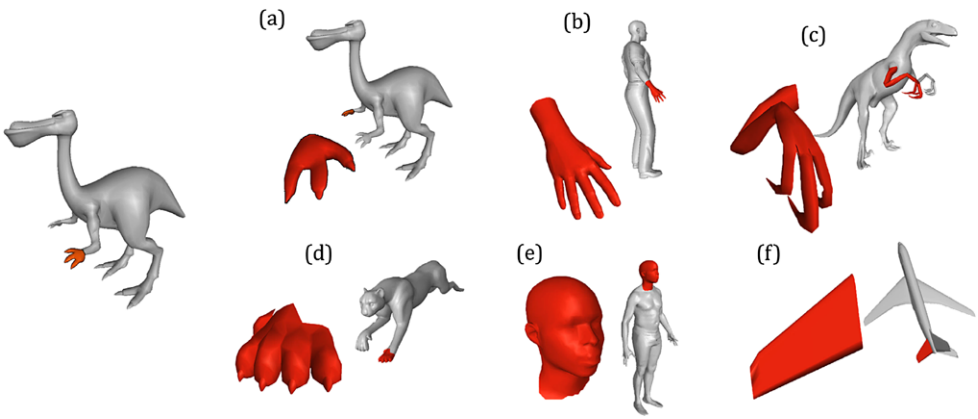

Fig. 13 Distance measure comparison. We measure the distance from hand of the dinopet to six other parts. Parts (**a**) through (**d**) are similar in spite of their large geometric variability, while parts (**e**) and (**f**) are not. The distance measurements are listed in Table 1

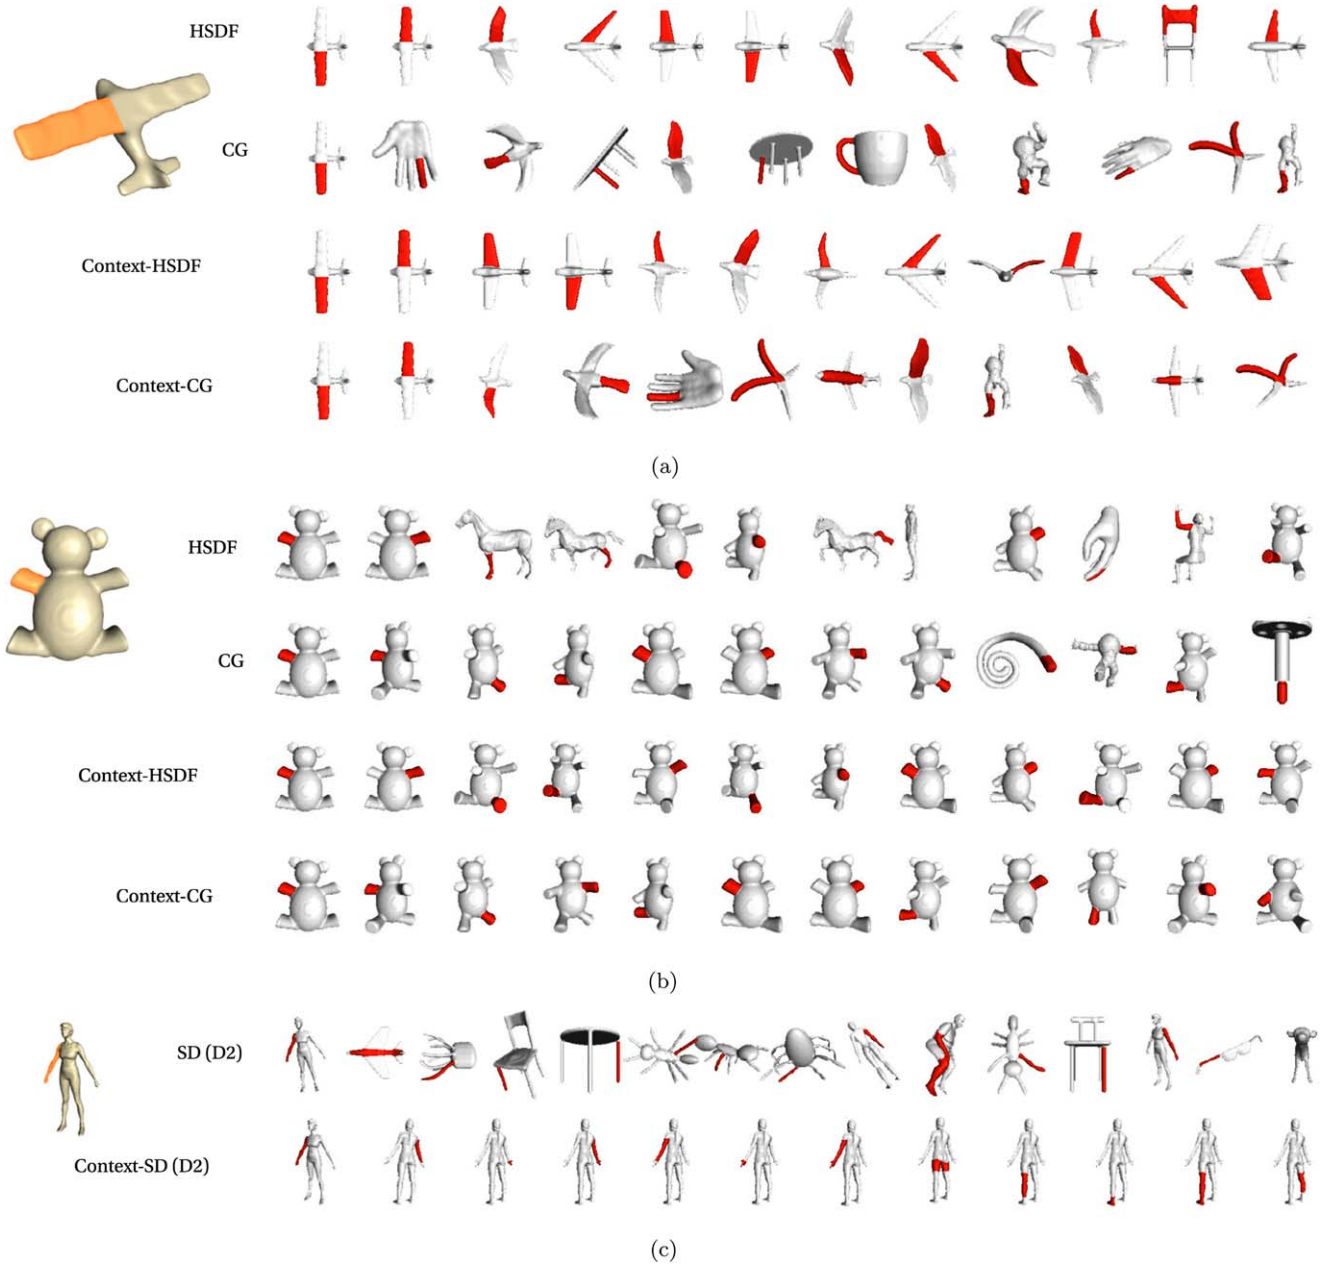

Fig. 14 Taking as example queries an airplane wing (**a**), the hand of a teddy bear (**b**) and the arm of a woman (**c**), we see significantly better query results when moving from a local distance measure to a context-aware one

ogous parts, enriching the database with semantic metadata.

5.1 3D Model Parts Retrieval

Using the distance measure defined in Sect. 4 we developed a simple part retrieval application. The user loads a model, which is automatically partitioned. The user can select a part p and search for similar parts in the database. The database models are segmented to parts, each retaining its partitioning hierarchy, and pre-calculated geometric attributes. We scan the database, and for each part p', calculate the contextual distance measure $D(p, p')$. We sort the results and return the top list of matching parts. Several example queries can be seen in Fig. 15.

We also allow to search for analogies in a set of models (Fig. 19). Given a source model, and k target models, we attempt to find a maximal correspondence between the source and each of the target models. This is done in a greedy fashion, which queries successively each part in the source model, over the subset of target models. The best match is

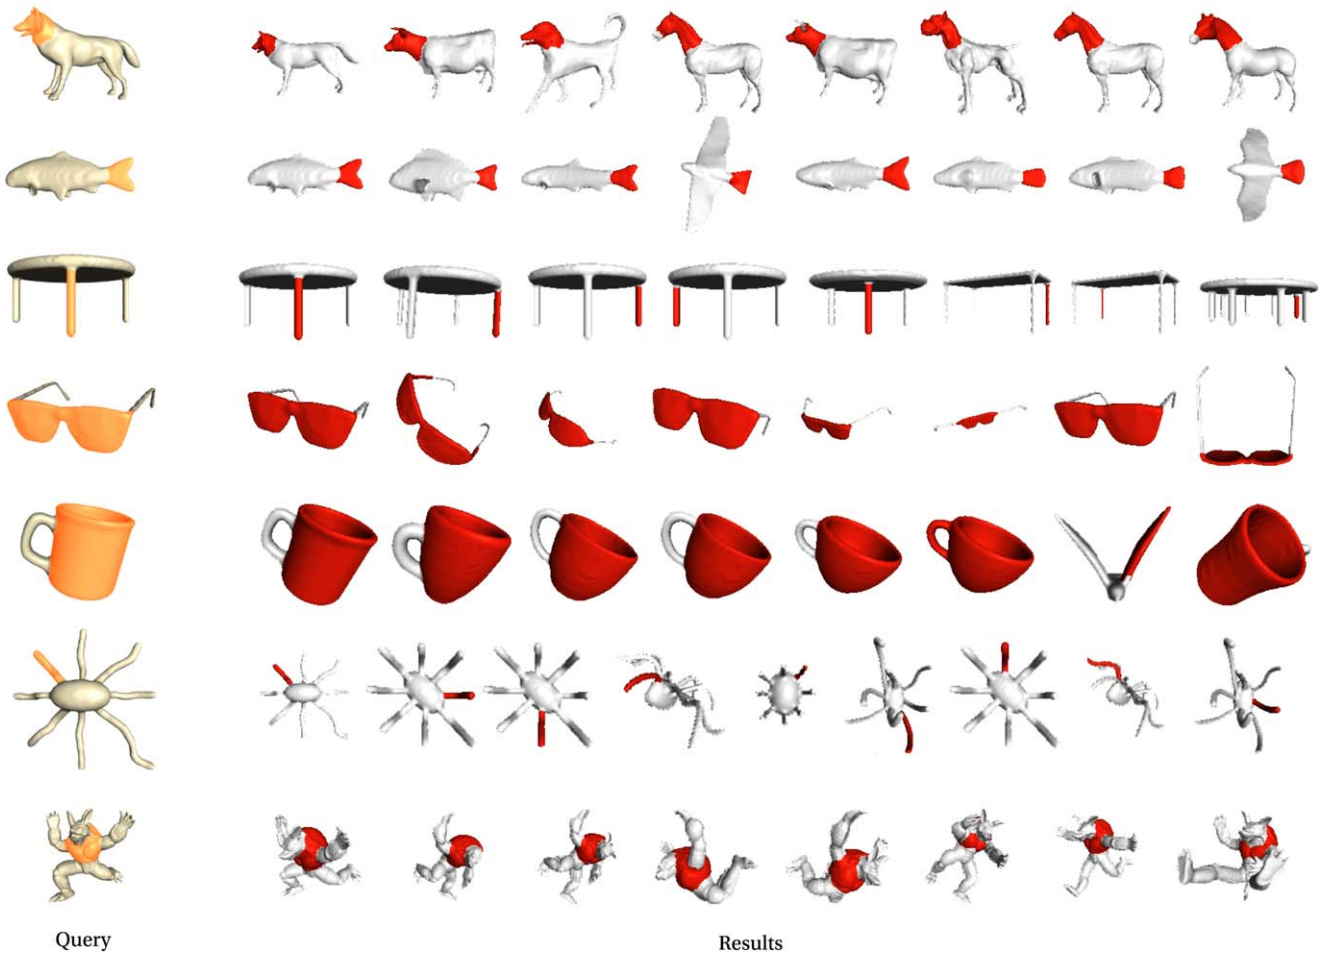

Fig. 15 Results of several part queries. On the left is the query part and on the right the search results. All examples are from the SHREC database

selected, and the matching continues on the parts adjacent to it.

We have conducted experiments on two databases. The first is the SHREC water-tight models database (Giorgi et al. 2007) which contains 400 models in various categories such as men, women, animals, ants, planes, chairs, tables etc. (for samples see Fig. 2). The second database contains models gathered from various sources, including the Princeton Shape Benchmark database (Shilane et al. 2004). This database contains 300 models.

For both databases we have partitioned the models using 4 iso-values, which results in up to four levels of partitioning. The partitioning resulted in 3562 distinct parts in the SHREC database, and 4711 parts in our second database. For an example of the partitioning results see Fig. 9.

Due to the fact that the parts of the models we use are not categorized, it has been difficult to quantize the success of our algorithm and compare it to other algorithms. To the best of our knowledge, no previous works have yet compiled a segmented semantic part database of 3D models, and conducted extensive testing on it. We hope our efforts will be the first step in establishing such a benchmark.

We have run queries for each part in the SHREC database and tested the results using the Nearest-neighbor test (Shilane et al. 2004). Using the context-aware distance measure we were able to achieve 97.7% accuracy, compared to 93% accuracy using the local HSDF measure. Additionally, we defined several test categories such as Airplane wing, Armadillo hand, Human leg, and ran queries on parts within these categories. Each query was run on local and context-aware distance measures. The context-aware distance measures outperformed the local distance measures in each query with no exceptions. A plot of the results can be seen in Fig. 16.

We illustrate the effectiveness of our technique on other classes of objects by adding around 30 CAD models to the database. We also include variations of the same model and search for a specific part. Figure 17 shows that although geometrically the parts may vary, since we use context-based measures, we find the correct analogous parts in other mod-

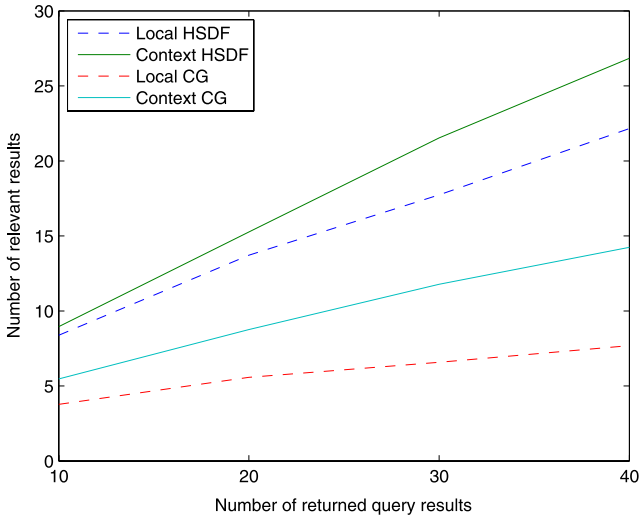

Fig. 16 We performed part queries in several categories such as Human leg, Armadillo leg and Airplane wing. We have compared the results using two local distance measures and their context-aware counterparts. The number of relevant results returned with the use of contextual analogies surpassed its local counterpart over all queries

els. Furthermore, we inserted into the parts database different partitioning of the same object, and used such parts as queries. Still, regardless of the partitioning, the results returned similar parts from all the copies of the model, illustrating a level of robustness to specific (possibly incorrect) partitioning of objects (Fig. 18).

All statistics were gathered on a 2.4 GHz dual core Windows XP machine. The pre-processing steps for building the SHREC database are summarized in Table 2 along with timing information. A query takes on average 600 ms to cover all parts in the database, and return the relevant results.

Table 2 Pre-processing times for the SHREC database

Task	Time (mins)	Comment
SDF Values	10	400 models
Auto Partitioning	5	400 models
Conformal Geometry Sig.	60	3652 parts
Shape Dist. Sig.	90	3652 parts

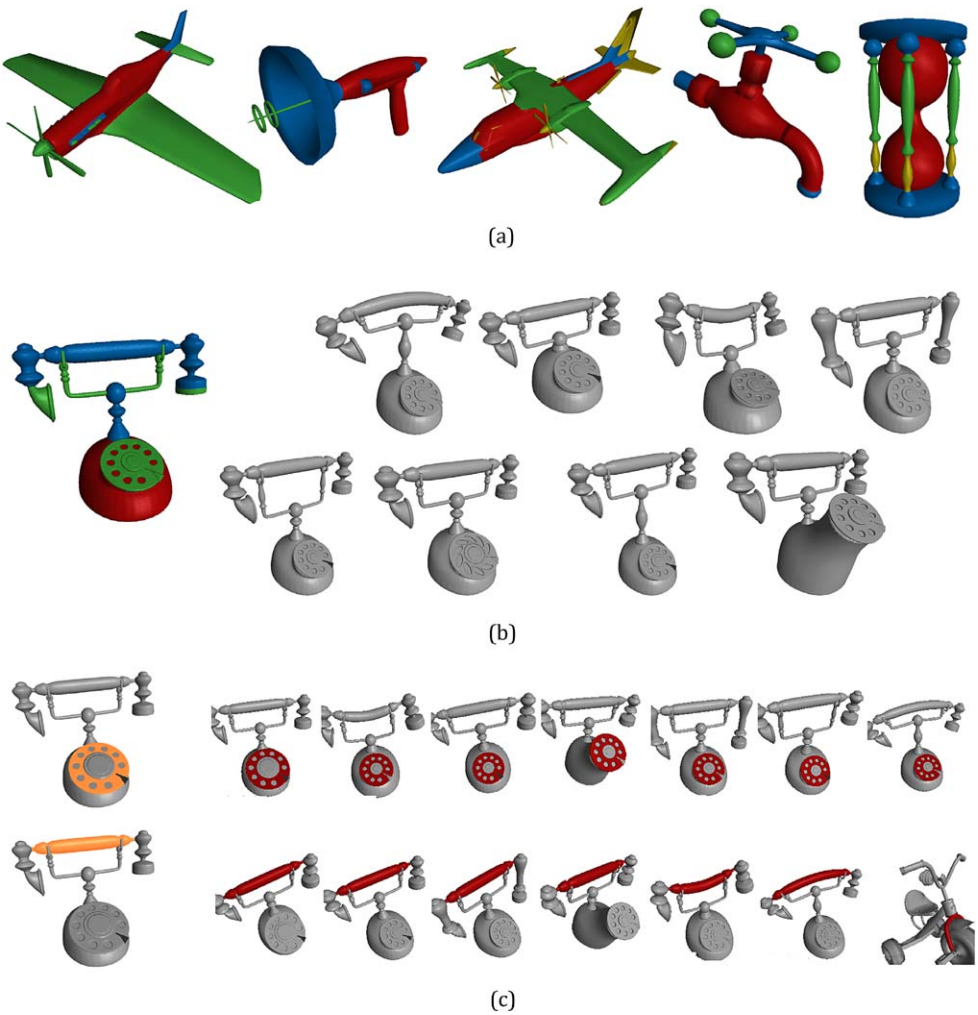

Fig. 17 CAD Models. (**a**) We added a class of mechanical models to the database and partitioned them. (**b**) We added several variations of the phone model created using iWires (Gal et al. 2009). (**c**) Querying for the dial and handle of the phone returns the correct parts in the different variations

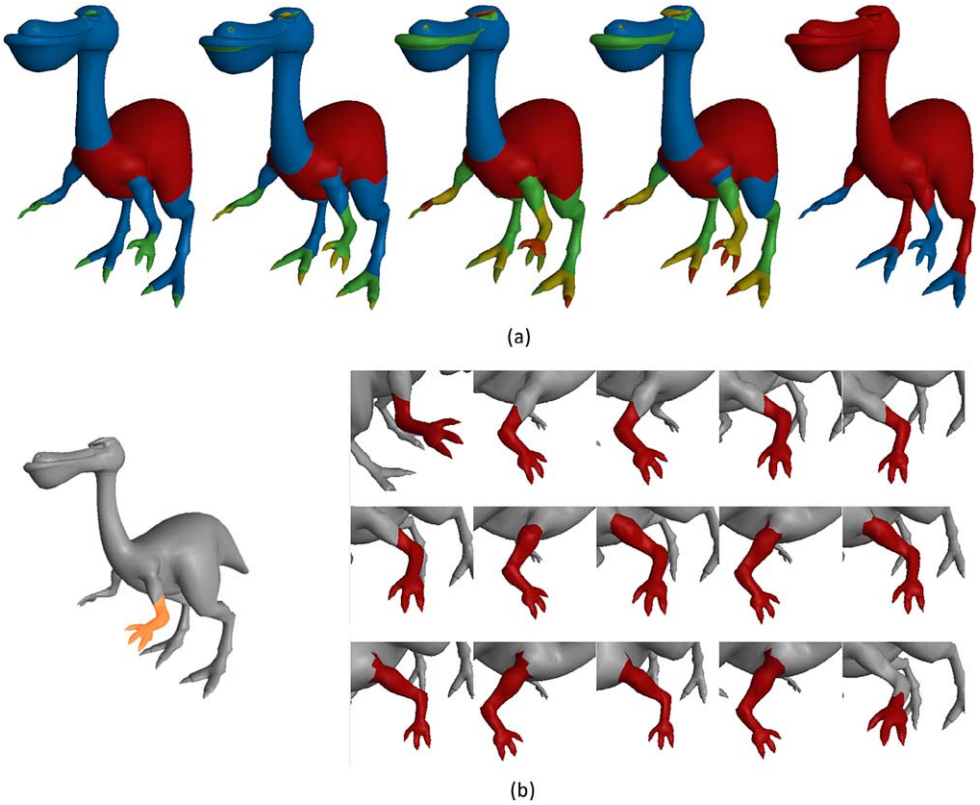

Fig. 18 Five distinct partitioning of the dinopet model are inserted into the database (**a**). Even though the part hierarchy is different for each model, querying the dinopet hand (**b**) returns matching hands from all dinopet variants

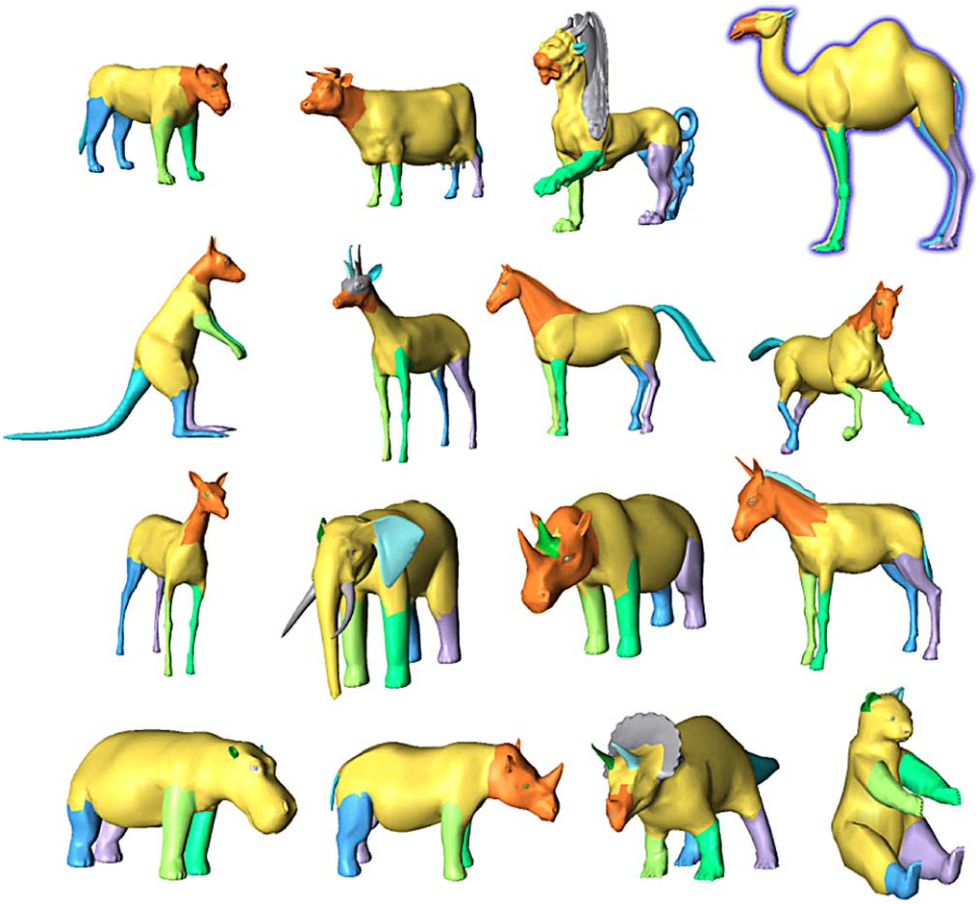

Fig. 19 Analogies between parts of whole objects, as indicated by matching colors

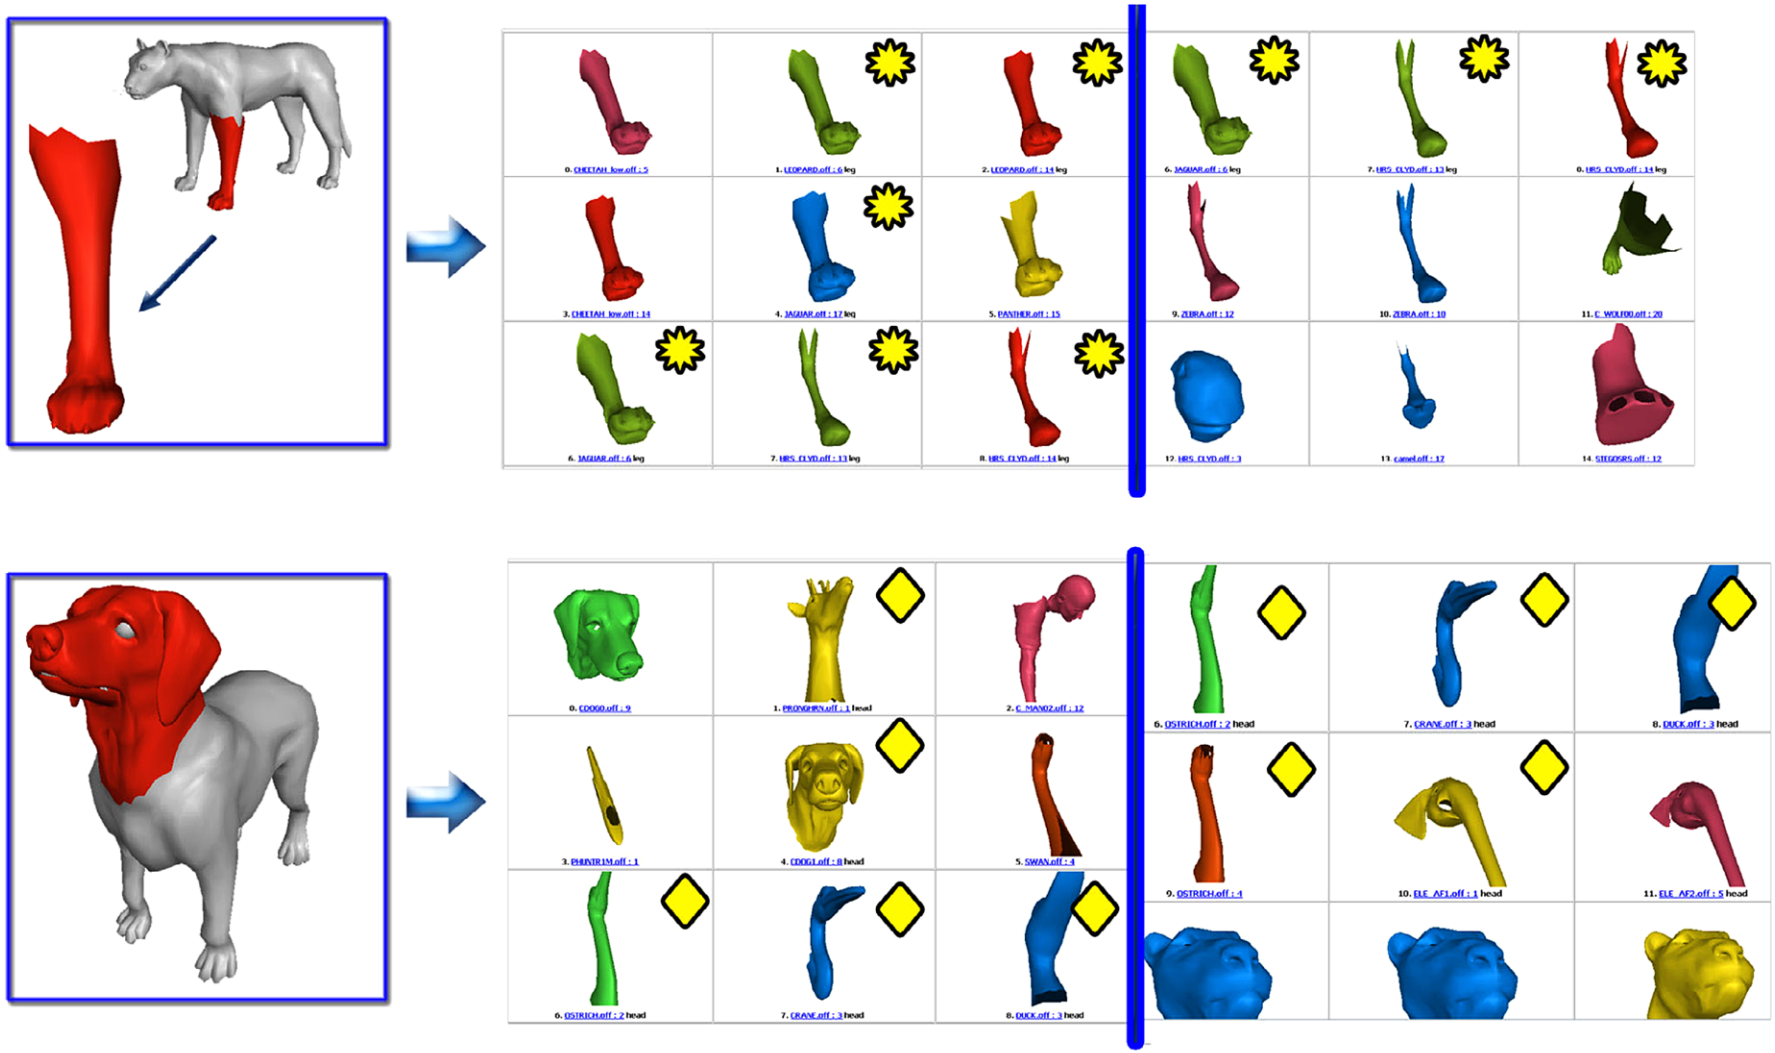

Fig. 20 (Color online) Two examples of automatic contextual tag transfer. *Top*: We search for the cheetah's leg. All results marked with a *yellow asterisk* have already been tagged as '**leg**'. *Bottom*: We search for the dog's head. All results marked with a *yellow diamond* have already been tagged '**head**'. These tags are now transferred to the query parts

5.2 Part Annotation

Using the contextual distance measure, we can now transfer user supplied annotations from one part to others in our database automatically. We developed a simple interface, in which a user may select a part (of any level in the hierarchy of the corresponding object) and annotate it with one or more textual tags. The tag is then associated with the part, and kept in the database.

Given a part p which we wish to automatically annotate, we define it as a query and search the database, retrieving a set of results. We discard all but the first 20 results from the set, denoting the resulting subset by R, and build a set of tags $T(R)$ containing all tags attached to parts in R. For each tag $t \in T(R)$ we define a tag importance measure:

$$C(t) = \sum_{r \in R_t} \frac{1}{D(p,r) - 1},$$

where $R_t = \{r \in R | t \in r\}$ and D is the shape context distance measure defined in Sect. 4.

We associate part p with all tags t such that $C(t) > m$, where the threshold m is set to 100 throughout. The tags are attached to the parts and saved in the database.

We allow the user to perform an annotation transfer on all tags found in the database, or only on selected tags. Consequently we can perform text searches in the database, searching for specific tags, such as "'ear"', "'head"', "'thin"', "'wide"' etc. (Fig. 20).

6 Conclusions and Future Work

We have presented a framework which automatically finds part analogies among 3D objects. The method first partitions a given 3D object to create a part hierarchy, and then defines a signature for each part. This signatures draws not only from the properties of the part itself, but from the relations between the part and the whole object. Using these signatures we defined an effective context-aware distance measure that can find analogous parts among other objects, which are not necessarily similar as a whole.

We have shown that such part analogies can support part search queries in a shape retrieval application. We have also used them to add semantic information to the objects by carrying information defined on one part (e.g. tags) to analogous parts in other objects.

The current method relies on the initial hierarchical partitioning of the objects. A stronger approach would attempt to analyze or partition the object in various ways depending on the query context. This would allow more flexible analogies to be found and better support to partial matching which is not restricted to the given partitioning. Such an investigation is left for future work. Other possible future directions include the use of contextual distance measures with different signatures and the extension of the tagging application to full semantic taxonomies of objects databases.

References

Alexa, M. (2002). Recent advances in mesh morphing. *Computer Graphics Forum*, *21*(2), 173–196.

Attene, M., Biasotti, S., Mortara, M., Patane, G., Spagnuolo, M., & Falcidieno, B. (2006a). Computational methods for understanding 3D shapes. *Computers & Graphics*, *30*(3), 323–333.

Attene, M., Falcidieno, B., & Spagnuolo, M. (2006b). Hierarchical mesh segmentation based on fitting primitives. *The Visual Computer*, *22*(3), 181–193.

Attene, M., Robbiano, F., Falcidieno, B., & Spagnuolo, M. (2007). Semantic annotation of 3D surface meshes based on feature characterization. In *Proceedings SAMT 2007. Lecture notes in computer science* (pp. 126–139). Berlin: Springer.

Belongie, S., Malik, J., & Puzicha, J. (2001). Matching shapes. In *Proceedings ICCV* (pp. 454–463).

Ben-Chen, M., & Gostman, C. (2008). Characterizing shape using conformal factors. In *Proceedings of Eurographics workshop on shape retrieval 2008*.

Bilmes, J. (1997). *A gentle tutorial on the em algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models*.

Boykov, Y., Veksler, O., & Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*(11), 1222–1239.

Bustos, B., Keim, D. A., Saupe, D., Schreck, T., & Vranić, D. V. (2005). Feature-based similarity search in 3D object databases. *ACM Computing Survey*, *37*(4), 345–387.

Camossi, E., Giannini, F., & Monti, M. (2007). Deriving functionality from 3D shapes: Ontology driven annotation and retrieval. *Computer-Aided Design & Applications*, *4*(6), 773–782.

Choi, H., Choi, S., & Moon, H. (1997). Mathematical theory of medial axis transform. *Pacific Journal of Mathematics*, *181*(1), 57–88.

Cornea, N.D., Demirci, M.F., Silver, D., Shokoufandeh, A., Dickinson, S., & Kantor, P.B. (2005). 3d object retrieval using many-to-many matching of curve skeletons. In *SMI '05: proceedings, shape modelling and applications conference*. Los Alamitos: IEEE Computer Society.

Funkhouser, T., & Shilane, P. (2006). Partial matching of 3D shapes with priority-driven search. In: *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing* (pp. 131–142). Eurographics Association, Aire-la-Ville, Switzerland.

Funkhouser, T., Kazhdan, M., Min, P., & Shilane, P. (2005). Shape-based retrieval and analysis of 3D models. *Communications of the ACM*, *48*(6), 58–64.

Gal, R., & Cohen-Or, D. (2006). Salient geometric features for partial shape matching and similarity. *ACM Transactions on Graphics*, *25*(1), 130–150.

Gal, R., Shamir, A., & Cohen-Or, D. (2007). Pose oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics*, *13*(2), 261–271.

Gal, R., Sorkine, O., Mitra, N., & Cohen-Or, D. (2009, to appear). iWIRES: An analyze-and-edit approach to shape manipulation. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, *28*(3).

Gelfand, N., Mitra, N. J., Guibas, L. J., & Pottmann, H. (2005). Robust global registration. In *Proc. symp. geom. processing*.

Giorgi, D., Biasotti, S., & Paraboschi, L. (2007). *Shape retrieval contest 2007: Watertight models track* (Tech. Rep). CNR-IMATI, Friedrich-Wilhelms-Universität Bonn.

Johnson, A., & Hebert, M. (1999). Using spin images for efficient object recognition in cluttered 3D scenes. *PAMI*, *21*(5), 433–449.

Katz, S., & Tal, A. (2003). Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Transactions on Graphics (Proceedings SIGGRAPH 2003)*, *22*(3), 954–961.

Kraevoy, V., & Sheffer, A. (2004). Cross-parameterization and compatible remeshing on 3D models. *ACM Transactions on Graphics*, *23*(3), 861–869.

Kullback, S., & Leibler, R. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, *22*(1), 79–86.

Novotni, M., Degener, P., & Klein, R. (2005). *Correspondence generation and matching of 3D shape subparts* (Tech. Rep). CG-2005-2, Friedrich-Wilhelms-Universität Bonn.

Osada, R., Funkhouser, T., Chazelle, B., & Dobkin, D. (2002). Shape distributions. *ACM Transactions on Graphics*, *21*(4), 807–832.

Schreiner, J., Asirvatham, A., Praun, E., & Hoppe, H. (2004). Inter-surface mapping. *ACM Transactions on Graphics*, *23*(3), 870–877.

Shamir, A. (2007, to appear). A survey on mesh segmentation techniques. *Computer Graphics Forum*.

Shapira, L., Shamir, A., & Cohen-Or, D. (2008). Consistent mesh partitioning and skeletonisation using the shape diameter function. *The Visual Computer*, *24*(4), 249–259.

Shilane, P., Min, P., Kazhdan, M., & Funkhouser, T. (2004). The Princeton shape benchmark. In *Shape modeling international*.

Sumner, R. W., & Popovic, J. (2004). Deformation transfer for triangle meshes. *ACM Transactions on Graphics*, *23*(3), 399–405.

Sundar, H., Silver, D., Gagvani, N., & Dickinson, S. (2003). Skeleton based shape matching and retrieval. In *SMI '03: Proceedings, shape modelling and applications conference* (p. 130). Los Alamitos: IEEE Computer Society.

Tangelder, J. W., & Veltkamp, R. C. (2008). A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications*, *39*(3), 441–471. doi:10.1007/s11042-007-0181-0.

Zabih, R., & Kolmogorov, V. (2004). Spatially coherent clustering using graph cuts. *CVPR'04*, *02*, 437–444.

Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., & Tagliasacchi, A. (2008). Deformation-drive shape correspondence. *Computer Graphics Forum*, *27*(5), 1431–1439. (Special Issue of Symposium on Geometry Processing 2008.)