

UNIVERSITÀ CA' FOSCARI DI VENEZIA
Dipartimento di Informatica
Technical Report Series in Computer Science

Rapporto di Ricerca CS-2009-8

Luglio 2009

A. Erdem, A. Torsello

A Game Theoretic Approach To Jointly Learn
Shape Categories and Contextual Similarities

Dipartimento di Informatica, Università Ca' Foscari di Venezia
Via Torino 155, 30172 Mestre–Venezia, Italy

A Game Theoretic Approach To Jointly Learn Shape Categories and Contextual Similarities

Aykut Erdem and Andrea Torsello

Dipartimento di Informatica
Università “Ca’ Foscari” di Venezia
{erdem,torsello}@dsi.unive.it

July 27, 2009

Abstract

Psychology studies from the late 70s showed that our perception of similarities is strongly influenced by the underlying category structure [21]. Further, the search of a model for representing and evaluating the similarities between shapes in a perceptually coherent way is still an open issue [12]. Starting from these observations, we aim at jointly learning the categories from examples and the similarity measures related to them. There is a chicken and egg dilemma here: class knowledge is required to determine perceived similarities, while the similarities are needed to extract class knowledge in an unsupervised way. In this paper, we present a game theoretic approach to compute 2D shape categories using disconnected skeletons [1]. The approach, based on [20], provides us with both the cluster information needed to extract the categories, and the relevance information needed to compute the category model and, thus, the similarities. The approach has been tested on a database of 1000 shapes and compared favorably with the pairwise clustering approaches working on similarities that do not make use of the underlying contextual information and a feature-based approach similar to ours.

1 Introduction

The unsupervised learning of shape categories is a central problem in computer vision with significant theoretical and practical impacts. There are two interrelated aspects to the problem: The first is the discovery of the shape categories present, and this can be effectively addressed as a problem of clustering shapes, while the second is the generalization of the class properties, i.e., the ability to assign each newly encountered shape to one of the extracted classes, or to recognize it as an outlier. Fundamental to both tasks is the problem of determining how similar two shapes really are.

These issues have been extensively studied with geometric characterizations of shape using both simple descriptors such as landmark points on the boundary [4], or more complex ones such as curve descriptors [8]. Shape-classes can

then be located by vectorizing the shape-attributes and applying standard central clustering techniques to the shape-vectors, while the problem of determining the membership to a class can be solved by performing principal components analysis on the class covariance matrix [4], or on a shape-manifold [8]. An alternative to the use of a single vectorial representation of the shape’s geometry is to use a structural abstraction where the object is divided into atomic components whose arrangement is then represented using a relational graph [7, 24]. Typically in this context, the similarity between two shapes is a measure of how well the primitives forming the shapes and/or their spatial organizations agree, and the assessment of whether a shape belongs to a particular class is performed by comparing in isolation the shape to one or more prototypes and by applying the nearest neighbor rule, while categories can be extracted using pairwise clustering [19].

One problem with these approaches is that they all assume the existence of a single *universal* measure of similarity between shapes, often requiring metric properties as well, while psychological experiments suggest that the human perception of similarity is not only non-metric [5], but also strongly dependent on the surrounding context [21, 12]. In particular, the observed variation within a shape-class is fundamental for determining the perception of the similarities of the shapes belonging to that class. Recently, this issue has also been surfaced from a computational point of view [23, 9].

Furthermore, there has been a growing interest in trying to characterize the modes of variation of shape abstracted in terms of graphs [18], which in turn induces a measure of the similarity of a shape to the whole class. However, it does nothing to determine the perceptual similarity to any particular shape belonging to the class, a problem that is central in *query by example* approaches in content based image retrieval systems. In [2], a new skeleton based shape retrieval system based on *disconnected skeleton* representation [1] was proposed. It provides a similarity measure that depends on the geometric characteristics of the shape categories present in the system. There, however, the category information had to be known a priori and was not inferred from the data.

In this paper, we propose a game theoretic approach to compute shape categories in an unsupervised way. There is a chicken and egg problem here: Class knowledge is required to determine perceived similarities, while the similarities are needed to extract class knowledge. We solve this problem using a EM-like approach where we iteratively estimate the class memberships and maximize for the parameters of our category representation. The expectation of class membership is obtained by adopting a game theoretic clustering framework presented in [20]. Then the similarities are computed as the edit-distance of a skeletal representation presented in [2] using the newly estimated cost coefficients. Central to the approach is the ability of the clustering framework to provide both the cluster information needed to extract the categories, and the relevance information (or the degree of membership) needed to compute the category model, and thus the similarities, in a robust way. Interestingly, the contextual similarity defined in [2] is not symmetric, making the ability of the game-theoretic approach to deal with asymmetric affinities particularly attractive.

2 Disconnected Skeletons and Category Influenced Matching

First introduced by Blum [3], skeleton based shape representations are one of the most common representation schemes for generic shape recognition, *e.g.* [24, 15, 13, 11], as they capture part structure and provide insensitivity to articulations and occlusions. However, in practice, two visually very similar shapes might have structurally different skeletons, hence this instability issue should be resolved either in extracting the skeleton or in the matching process. In this regard, *disconnected skeletons* [1] provide an alternative solution as the method aims at obtaining a coarse yet a very stable skeleton representation from scratch.

Disconnected skeletons are defined in terms of a special distance surface (Aslan surface), the level curves of which are increasingly smoothed versions of the initial shape boundary, and which has a single extremum point that captures the center of a blob-like representation of the shape (Fig. 1(a)). Each branch extracted from this surface is classified as either *positive* or *negative*, identifying whether it originates from a positive curvature maxima (a protrusion) or a negative curvature minima (an indentation). Among the extracted branches, at least two positive and two negative branches reach the shape center, and these are called *major branches* since they represent the most prominent visual features of the shape. All the other branches terminate at some *disconnection points* where a positive branch and a negative branch collide (Fig. 1(b)). It has been shown that these points are very stable under visual transformations such as articulation and deformation of parts.

The termination concept is artificially extended to the major positive branches and the relative organization of the branches are expressed in terms of the location of their disconnection points on a coordinate frame based on the major negative branches (Fig. 1(c)). Note that the coordinate frame can be constructed in a number of ways depending on the choice of major branch, thus all possible descriptions are stored. The skeletal attributes used to represent each skeleton branch are simply its type, the location of its disconnection points (r, θ) , and its length l measured in the formed coordinate frame.




Figure 1: Disconnected skeletons. (a) Level curves of Aslan surface (b) Positive and negative skeleton branches, respective drawn in blue and red (before pruning) (c) Spatial organization of skeleton branches (taken from [2]).

In [2], disconnected skeletons are represented as *rooted attributed depth-1 trees* (Fig. 2) and tree-edit distance is used to match these structures. Moreover, an ordering scheme based on the major negative branches is introduced in order to speed up the matching process. In a tree-edit distance based algorithm, a critical issue is how the edit operations are defined. The cost functions not only affect the success of the algorithm but also its overall time complexity. In [2], Baseski *et al.* used the context information about the category of one of the shapes to be matched to determine the cost functions, rendering the asymmetric role of the shapes in comparison and thus the distance function itself. The cost functions are computed on the basis of category specific statistics about the skeletal attributes that are stored in an auxiliary tree union structure. In this version, the cost function for the label **change** operation is defined in terms of a generic cost function. The idea resembles Mahalanobis distance in that when the distance within the observed range of skeletal attributes, but rapidly increases outside of that region.




Figure 2: Tree abstractions of some disconnected skeletons (taken from [2]).

3 Grouping Game

In [20], a novel framework for grouping and clustering was presented which was derived from a game-theoretic formalization of the competition between the hypotheses of group membership. The basic idea is as follows: Let the hypotheses that each element belongs to a group compete with one-another, each obtaining support from compatible elements and competitive pressure from all the others. Competition will reduce the population assuming hypotheses that do not receive strong support from the rest, while it will allow populations assuming hypotheses with strong support to thrive. Eventually, all inconsistent hypotheses will be driven to extinction, while all the surviving hypotheses will reach an equilibrium with all receiving the same average support. Clustering was thus formalized as a repeated non-cooperative game where competition for class membership selects elements belonging to a coherent cluster.

Specifically, let $O = \{1, \dots, n\}$ be the set of available elements, for each pair of strategies $i, j \in O$, a_{ij} represents the payoff of an individual playing strategy i against an opponent playing strategy j . A *mixed strategy* is a probability distribution $\mathbf{x} = (x_1, \dots, x_n)^T$ over the available strategies O .

$$\Delta = \{\mathbf{x} \in \mathbb{R}^n : x_i \geq 0 \text{ for all } i \in O, \mathbf{1}^T \mathbf{x} = 1\},$$

where $\mathbf{1} = (1, \dots, 1)^T$, while the *support* of a mixed strategy $\mathbf{x} \in \Delta$, denoted by $\sigma(\mathbf{x})$, is defined as the set of elements chosen with non-zero probability: $\sigma(\mathbf{x}) = \{i \in O \mid x_i > 0\}$.

The expected payoff received by a player choosing element i when playing against a player adopting a mixed strategy \mathbf{x} is $(A\mathbf{x})_i = \sum_j a_{ij}x_j$, hence the expected payoff received by adopting the mixed strategy \mathbf{y} against \mathbf{x} is $\mathbf{y}^T A\mathbf{x}$.

The *best replies* against mixed strategy \mathbf{x} is the set of mixed strategies

$$\beta(\mathbf{x}) = \{\mathbf{y} \in \Delta \mid \mathbf{y}^T A\mathbf{x} = \max_{\mathbf{z}} (\mathbf{z}^T A\mathbf{x})\}.$$

A strategy \mathbf{x} is said to be a *Nash equilibrium* if it is the best reply to itself, i.e.,

$$\forall \mathbf{y} \in \Delta \quad \mathbf{x}^T A\mathbf{x} \geq \mathbf{y}^T A\mathbf{x}. \quad (1)$$

Nash equilibria abstracts the main characteristics of a group: internal homogeneity, that is, a high mutual support of all elements within the group, and external dishomogeneity, or low support from elements of the group to elements that do not belong to the group. In fact, any element $i \in \sigma(\mathbf{x})$ of a Nash equilibrium \mathbf{x} receive from \mathbf{x} the same expected payoff $(A\mathbf{x})_i = \mathbf{x}^T A\mathbf{x}$, while elements not in $\sigma(\mathbf{x})$ receive a lower or equal support from the elements of the group. In this sense the elements belonging to a cluster are those in the support stable equilibria, and these are found using the replicator dynamics [22], a well-known formalization of a natural selection process.

The main characteristics of the framework are that it is generic, as it can deal with asymmetric as well as negative affinities; it does not require *a priori* knowledge of the number of clusters as it is inherently a multi-figure/ground discrimination process; and it provides immediate measures of both the cohesiveness of the cluster in the form of its average payoff $\mathbf{x}^T A\mathbf{x}$, and of the participation of an element to the cluster. In fact the value x_i can be interpreted as a degree of participation of element i to the cluster defined by the stable point \mathbf{x} .

4 The Proposed Method

In this study, we attempt to solve the interrelated problems of discovering shape categories and computing the corresponding contextual similarities using a EM-like approach where we iteratively estimate the class memberships and maximize for the parameters of our category representation. The expectation of class membership is obtained by adopting the game theoretic clustering framework summarized in Section 3. Then the similarities are computed as the edit-distance of a skeletal representation presented in [2] using the newly estimated cost coefficients. The details of these steps are as follows.

4.1 Discovering Shape Categories

Following [20], we define the shape category in terms of a clustering game where shapes present in the training set compete for category membership. The outcome of the competition is determined by the payoff or utility matrix $A = (a_{ij})$ which represents the gain a player obtains when playing the strategy (shape class) i against an opponent playing j . Initially, these payoffs simply correspond to the similarities among the given set of shapes obtained with $a_{ij} = \exp\left(-\frac{(\text{dist}(i,j))^2}{\sigma^2}\right)$ where σ is a scaling factor, and $\text{dist}(i,j)$ is the tree-edit distance between the disconnected skeletons of the shapes i and j .

Since no category information is available in the beginning, the initial similarities were computed in isolation without any context, thus A is a *symmetric* matrix. However, in the subsequent iterations, the category structure discovered in the previous step influences the similarity computations by differentiating the roles of the shapes in comparison. Now, each row index corresponds to a query shape whereas each column index is a shape which has a category label assigned by the previous grouping game (if it is not found to be an outlier), and the cost functions are determined by the context about the category of the second shape. Thus, this results in an *asymmetric* similarity matrix.

Given the payoff matrix A , we extract shape categories by applying a *peel-off* strategy. At first, we start with a grouping game that considers all the shapes and we extract a cluster by running the replicator dynamics. Following that, we define a new game on the set of remaining (unlabeled) shapes and reiterate the procedure until all directed-dominant sets are extracted. The game theoretic framework also provides us a direct way to evaluate the coherency of extracted clusters. Let $S \in \mathcal{S}$ be an extracted dominant set, the coherency of S can be computed as its average payoff $\mathbf{x}_S^T A \mathbf{x}_S \in [0, 1]$. By inspecting these values, we obtain an initial set $\mathcal{C} (\subseteq \mathcal{S})$ of coherent shape categories which is formed by the clusters $S \in \mathcal{S}$ with $\mathbf{x}_S^T A \mathbf{x}_S > \zeta_1$. This allows us to discard incoherent classes hence enforcing robustness in the extraction process.

The payoff information can also be used to assign additional members to the clusters in \mathcal{C} . To compute the similarity between a shape i to a cluster S , we use the weighted similarity function $\gamma_S(i) = \frac{(A \mathbf{x}_S)_i}{\mathbf{x}_S^T A \mathbf{x}_S}$. We evaluate this similarity measure for every unlabeled shape i and assign it to the most similar cluster if $\gamma_S(i) \geq \zeta_2$. Otherwise, it is considered as an outlier shape which does not belong to any of the extracted categories. The ability of assigning an unclustered object either to a category or to the *outlier* class is instrumental to the generalization capabilities. Note that the outlier class should be interpreted as a “don’t know” label where the approach cannot say anything about the shape rather than recognizing the shape as a new class not seen in the other examples.

After reassigning the leftover elements, we re-examine the dominant sets that were rejected by the first thresholding step and check whether they became more coherent with the removal of the reassigned elements. To evaluate their coherency we use an hysteresis strategy: we accept the groups with $|S| > 3$ whose average payoffs $\mathbf{x}_S^T A \mathbf{x}_S > \zeta_3$, with $\zeta_3 < \zeta_1$. The purpose of this hysteresis is to reduce the effect the implicit change in scale induced by the peel-off strategy and to increase robustness with respect to the scaling factor σ .

4.2 Computing Contextual Similarities

To model the influence of the discovered category structure on the computation of shape similarities, we adopt the tree-edit distance based shape matching method proposed by Baseski *et al.* [2]. Here, unlike the supervised approach in [2], we form the union based on the disconnected skeleton trees of the shapes that are clustered together by the game-theoretic approach. Further, in the computation of the label **change** cost function, we substitute the minimum and maximum values of the skeletal attributes in the category with soft bounds that make use of the membership information supplied by the clustering framework. In particular, we use the weighted mean $\mu_{\mathbf{x}}$ and weighted standard deviation $\sigma_{\mathbf{x}}$ (Eqn. 2) to determine the range $\mu_{\mathbf{x}} \pm 3\sigma_{\mathbf{x}}$ which has experimentally shown to account for the shape variability and provide a robust inference process.

$$\mu_{\mathbf{x}} = \frac{\sum_{i=1}^n x_i y_i}{\sum_{i=1}^n x_i}, \quad \sigma_{\mathbf{x}} = \sqrt{\frac{1}{1 - \sum_{i=1}^n x_i^2} \sum_{i=1}^n x_i (y_i - \mu_{\mathbf{x}})^2} \quad (2)$$

In obtaining the affinity matrix $A = (a_{ij})$ at time step $t > 0$, we introduce a soft indexing scheme where we propagate the information about the similarities to the extracted classes: When computing the similarity between the query shape i and the shape j , if j belongs to a cluster S extracted in the previous step, we multiply the similarity influenced by the new category information, with the similarity of the shape i to the cluster S normalized with respect to the most similar category. This allows us to bias the similarities towards the previously extracted clusters, thus propagating the membership throughout the iterations. Clearly, if j is an outlier shape, *i.e.* no category information is available about it, we keep the original distance which does not utilize any context information. Moreover, the corresponding multiplier b_{ij} is taken as 1.

$$a_{ij} = b_{ij} \times \exp(-(dist(i, j))^2 / \sigma^2) \quad (3)$$

$$\text{where } b_{ij} = \begin{cases} 1 & \text{if } j \text{ is an outlier} \\ \frac{\gamma_S(i)}{\max_{T \in \mathcal{C}} \gamma_T(i)} & \text{if } j \in S \end{cases}$$

4.3 The Stopping Criteria

For the convergence of the proposed approach, different stopping criteria can be defined. From the strongest to the weakest, these are: (1) the time when both the resulting clusters and the membership values of each elements in each cluster do not change, (2) the time when the resulting clusters are the same as the previous ones, (3) the time when the change in the ratio of unlabeled (outlier) shapes to the number of all shapes is smaller than or equal to a specified threshold (ζ_4). In this work, we used the weakest criteria (earliest stopping time) since the resulting group and distance information, as well as the query performance appear to be relatively stable after meeting that condition.

4.4 Computational Complexity

The computational complexity of the proposed method can be characterized by analyzing the complexities of the underlying shape matching and clustering frameworks: The time complexity of the tree-edit algorithm is $\mathcal{O}(mn)$ where m and n denote the number of leaf nodes in tree abstractions of the shapes in comparison. The time complexity of finding a Nash equilibrium using the replicator dynamics is $\mathcal{O}(kn^2)$ where n is the number of objects (strategies) and k is the number of needed iterations.

5 Experimental Results

In order to evaluate the performance of the proposed approach, we used the shape database provided in [2] which contains a total of 1000 shapes from 50 shape categories, each having 20 examples. We start with obtaining the disconnected skeleton descriptions, where we additionally compute and store the radius of the maximal circles at the termination points which are normalized with respect to the radius of the maximal circle at the shape center. After the descriptions are formed, we iteratively run the proposed method with the empirically set parameters $\sigma^2 = 24$, $\zeta_1 = 0.85$, $\zeta_2 = 0.95$ and $\zeta_3 = 0.75$, and stop when $\zeta_4 \leq 0.005$. In this setting, the algorithm converges at the 2^{nd} iteration.

The initial (symmetric) affinities and the (asymmetric) contextual similarities learned from the identified category structure at the 1^{st} and 2^{nd} iterations are shown in Fig. 3(a)-(c), respectively. Here, elements that belong to the same semantic category have consecutive indexes. The shape categories extracted from these pairwise similarities are given in Table 2 and Table 3, respectively, where we show the shape with the highest membership score from the semantic category most represented in each class. Moreover, Fig. 4 shows the total entropy of the clustering results at each iteration, which is the sum of the entropies of the population distributions linked with each cluster. Once a class is first extracted, the membership distribution is tightly peaked around a few shapes that belong to the class. As the class parameters and thus the contextual similarities are improved, the confidence on the membership of other elements

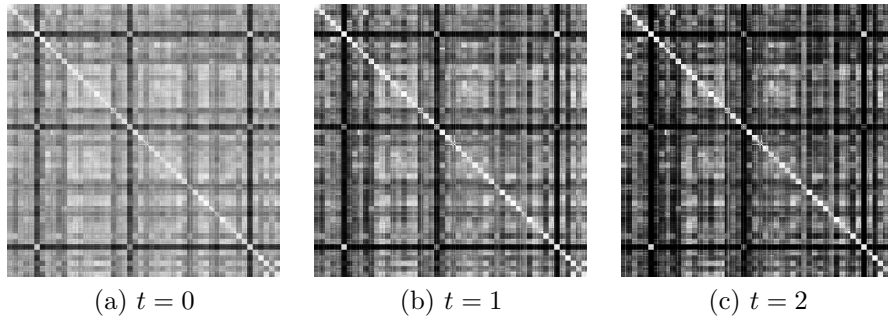


Figure 3: The evolution of similarities throughout the iterations. Observe that the suggested iterative procedure refines the initial (symmetric) similarities in the light of the discovered category structure, giving a better within versus between group separation, and resulting in asymmetric similarities.

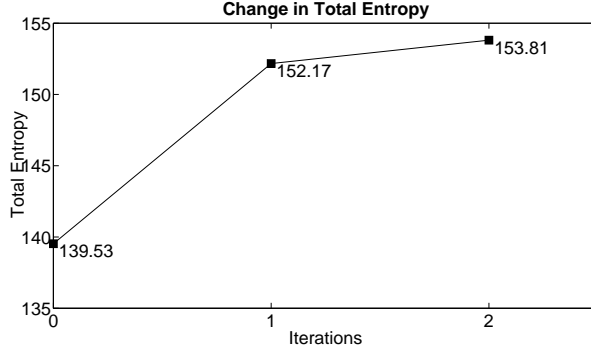


Figure 4: Total entropy throughout out each iteration.

to the class increases, yielding a flatter distribution and thus a higher entropy (Fig. 5). This is the normal behavior as long as the class membership does not change. On the other hand, if the new information changed the membership by eliminating an element from the class, then we would see a significant decrease in the entropy.

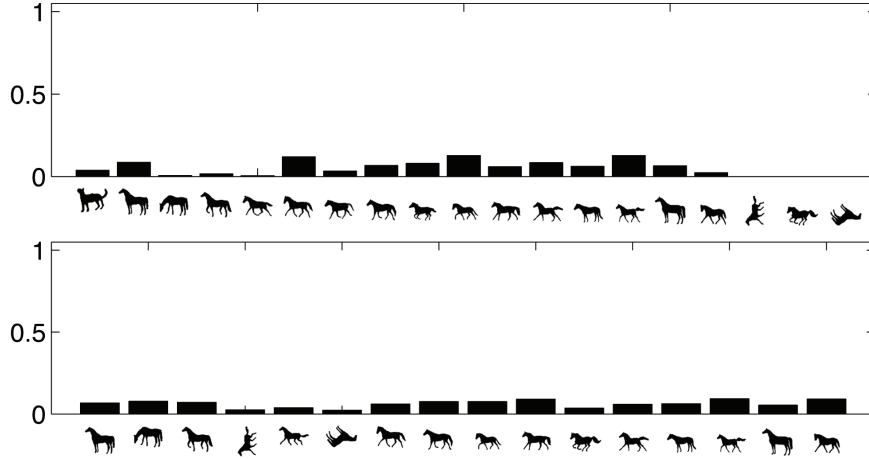


Figure 5: The initial (top row) and final (bottom row) population distributions of a discovered shape category. Total entropy increases from 2.55 to 2.71.

Table 1 shows some cluster validity measures on the classes extracted with our approach. The first measure is the standard Rand index [6], *i.e.* the ratio of agreements over all possible pairs. The second measure is a corrected version of the Rand index where the disagreements in the outlier class are not penalized, as this class is not supposed to form a coherent group. Note that the latter form of the Rand index favors more conservative approaches, where we prefer the approach not to label a shape when in doubt, while the former version favors bolder assignments where we prefer to make a few mistakes rather than not assign a shape to a class. Which version is to be preferred is clearly dependent on the application. The last measure is the normalized mutual information (NMI) [17] which measures the closeness between the class distributions and

The Method	Rand Index	Corrected Rand Index	NMI
Our method at $t=0$	0.9818	0.9929	0.8517
Our method (asymmetric case)	0.9854	0.9933	0.8722
Our method (symmetric case)	0.9885	0.9914	0.8780
Normalized Cut [16] (with # of classes=51)	0.9832	0.9833	0.8381
Normalized Cut [16] (with # of classes=61)	0.9848	0.9854	0.8380
Foreground Focus [9] (with # of classes=50)	0.9748		0.7329

Table 1: The quantitative evaluation of the clustering results.

the ground truth.

In an attempt to assess information content in the asymmetry of the similarity matrix, we also perform the same experiment using the same parameters but rendering the affinities symmetric before applying pairwise clustering. In this case, the approach requires 3 iterations to converge. The shape categories obtained in this experiment are given in Table 4. When the number of outlier shapes and the average precision recall values (Fig. 6) are considered alone, the symmetric case seems to work better than the asymmetric case. However, the difference between the plain and corrected Rand index show that the asymmetric approach is more conservative, i.e. it has a higher tendency to label shapes as belonging to an unknown class, but makes fewer misclassifications when it does assign shapes to a class, on the other hand the symmetric approach is more likely to assign shapes to a class, even when this results in more misclassifications.

We compared the results with several alternatives. The first, which should be seen as a baseline comparison, is performed by applying a pairwise clustering approach in order to extract the class structure, while assuming global, non-contextual similarities. Here we used Normalized Cut [16] as a baseline pairwise clustering approach. Note that the normalized cut approach requires the number of classes to be known *ab initio*. Here we choose two different values: 51 (the existing 50 semantic categories plus 1 for the outliers) and 61 (a number closer to the number of categories extracted with our approach). The additional number of classes is due to the fact that there can be a substantial semantic gap between appearance and categories, and allowing more freedom can result in better overall categorization. Indeed, as it can be seen in Table 1, normalized cut performs better with more degrees of freedom, but still performs significantly worse than the proposed approach.

The second approach we are comparing against is *Foreground Focus* [9]. This is an unsupervised algorithm proposed to learn categories from sets of partially matching image features. Just like our approach, it utilizes an EM-like algorithm to infer the categories. However, its goal is to learn relevant features rather than the actual contextual similarities. In order to compare with this method, we first form *Inner-Distance Shape Context* [10] descriptions of each shape by uniformly sampling 100 landmark points across the shape boundary and using a total of 5 inner-distance bins and 12 inner-angle bins. Earth Mover’s Distance (EMD) algorithm [14] is then used to compute the matchings of shape features

and similarities, and Normalized Cut [16] is used to determine the clusters. Here, the total number of extracted clusters is kept fixed at 50 (the actual number of shape categories exist in the database). Table 1 and Fig. 6 show that the performance of this approach is even significantly lower than the baseline normalized cuts over the skeletal distance. The huge difference can probably be explained by the lower descriptive power of the Inner-Distance Shape Context features with respect to disconnected skeletons, or bad performance of EMD matching algorithm.

The last comparison is with the label propagation method [23] and is limited to the retrieval performance of the contextual similarities. This method has three parameters which are used to construct the affinity matrix, the neighborhood size and the window size that are respectively set as $C = 0.275$, neighborhood size $K = 10$, window size $W = 250 \times 250$. When applied to the initial (symmetric) similarities, the approach offers a slightly better precision/recall (Fig. 6). However, note that the approach solves a slightly different problem; it concentrates only on improving retrieval rate and does not provide any category structure or an estimation of perceptually relevant similarities.




Figure 6: Average precision-recall curves.

6 Summary and Conclusion

In this paper, we presented an approach for the simultaneous discovering of 2D shape categories and the corresponding contextual similarities. This was achieved by adopting the game theoretic clustering approach introduced in [20] and by modifying the shape retrieval system presented in [2] in order to account for the uncertainty in the category information. The game theoretic framework naturally provides us the membership information about the extracted categories which quantifies this uncertainty, and is capable of dealing with the asymmetric similarities obtained using the contextual information. We

have demonstrated the potential of the proposed framework on a large shape database composed of highly varying 1000 shapes from 50 categories.






























































	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
															
Precision	1.00	0.80	0.90	1.00	1.00	1.00	1.00	1.00	0.78	1.00	0.74	1.00	0.94	0.82	0.94
Recall	0.95	0.60	0.90	0.90	0.85	0.95	0.90	0.90	0.35	0.50	0.85	0.90	0.75	0.90	0.85
Payoff	0.92	0.90	0.85	0.87	0.88	0.89	0.93	0.95	0.85	0.82	0.90	0.88	0.91	0.93	0.90
Entropy	2.63	2.43	2.30	2.33	2.40	2.81	2.59	2.83	1.69	2.18	2.73	2.28	2.56	2.58	2.40
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
															
Precision	0.92	0.53	1.00	0.94	0.93	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.95	1.00
Recall	0.60	0.45	1.00	0.80	0.70	0.40	0.25	0.85	0.45	0.55	0.25	0.45	0.25	0.90	0.70
Payoff	0.91	0.89	0.95	0.88	0.91	0.88	0.76	0.85	0.90	0.87	0.83	0.82	0.78	0.89	0.91
Entropy	2.33	2.44	2.91	2.27	2.47	1.94	1.29	2.39	1.96	2.15	1.60	2.06	1.44	2.55	2.47
	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45
															
Precision	0.43	1.00	1.00	1.00	0.54	0.60	0.76	1.00	1.00	0.91	1.00	1.00	0.83	1.00	1.00
Recall	0.30	0.45	0.25	0.90	0.70	0.30	0.65	0.95	0.45	0.50	1.00	0.85	0.25	0.95	0.85
Payoff	0.86	0.82	0.80	0.88	0.90	0.89	0.87	0.92	0.89	0.86	0.95	0.88	0.76	0.89	0.95
Entropy	1.85	2.04	1.46	2.44	2.87	1.87	2.30	2.67	2.14	2.13	2.80	2.38	1.45	2.36	2.71
	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
															
Precision	1.00	1.00	1.00	0.95	0.65	1.00	1.00	0.90	1.00	0.72	1.00	1.00	0.52	0.32	1.00
Recall	0.95	0.75	0.30	1.00	0.55	0.75	0.95	0.45	0.95	0.65	0.20	0.90	0.60	0.40	0.35
Payoff	0.96	0.88	0.76	0.94	0.90	0.91	0.90	0.86	0.89	0.89	0.81	0.93	0.93	0.87	0.84
Entropy	2.86	2.05	1.64	2.83	2.58	2.66	2.60	2.00	2.37	2.22	1.26	2.75	2.93	2.55	1.90
	61														
															
Precision	1.00														
Recall	0.35														
Payoff	0.83														
Entropy	1.87														

Table 2: The shape categories extracted from the initial (symmetric) affinities. The number of outlier shapes is 96.

References

- [1] C. Aslan, A. Erdem, E. Erdem, and S. Tari. Disconnected skeleton: Shape at its absolute scale. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(12):2188–2203, 2008.
- [2] E. Baseski, A. Erdem, and S. Tari. Dissimilarity between two skeletal trees in a context. *Pattern Recognition*, 42(3):370–385, 2009.
- [3] H. Blum. Biological shape and visual science. In *Journal of Theoretical Biology*, volume 38, pages 205–287, 1973.
- [4] T. F. Cootes, C. J. Taylor, and D. H. Cooper. Active shape models - their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, 1995.
- [5] D. W. Jacobs, D. Weinshall, and Y. Gdalyahu. Classification with non-metric distances: Image retrieval and class representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(6):583–600, 2000.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Precision	1.00	0.67	1.00	0.95	1.00	0.95	1.00	1.00	0.87	0.70	0.67	1.00	1.00	0.91	0.95
Recall	0.90	0.70	0.85	0.90	0.90	0.95	0.90	0.90	0.65	0.35	0.60	0.95	0.80	1.00	1.00
Payoff	0.96	0.95	0.92	0.94	0.95	0.94	0.96	0.97	0.91	0.88	0.94	0.94	0.94	0.94	0.92
Entropy	2.75	2.77	2.69	2.72	2.81	2.77	2.77	2.89	2.47	1.60	2.75	2.63	2.39	2.94	2.87
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Precision	1.00	0.50	0.75	1.00	0.94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Recall	0.65	0.35	0.15	1.00	0.75	0.70	0.45	0.20	0.90	0.60	0.40	0.45	0.25	0.25	0.80
Payoff	0.94	0.94	0.77	0.96	0.93	0.93	0.92	0.84	0.94	0.93	0.93	0.87	0.87	0.80	0.93
Entropy	2.39	2.55	1.34	2.92	2.59	2.46	2.03	1.38	2.78	2.44	2.03	2.13	1.61	1.57	2.71
	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45
Precision	1.00	0.62	1.00	0.83	1.00	0.37	0.86	0.92	1.00	1.00	0.91	0.95	1.00	1.00	1.00
Recall	0.50	0.50	0.55	0.25	0.80	0.35	0.60	0.60	0.95	0.50	0.50	0.95	0.90	0.25	1.00
Payoff	0.91	0.87	0.92	0.89	0.94	0.94	0.92	0.93	0.95	0.94	0.92	0.91	0.93	0.80	0.95
Entropy	2.10	2.19	2.28	1.61	2.68	2.75	1.94	2.43	2.79	2.30	2.38	2.72	2.78	1.58	2.77
	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
Precision	1.00	0.75	1.00	1.00	1.00	1.00	0.74	1.00	1.00	1.00	0.95	0.94	1.00	1.00	0.55
Recall	0.80	0.15	0.95	0.90	0.35	1.00	0.70	0.75	0.90	0.60	0.95	0.75	0.20	0.85	0.55
Payoff	0.90	0.83	0.96	0.95	0.84	0.97	0.95	0.95	0.95	0.91	0.93	0.93	0.86	0.96	0.94
Entropy	2.56	1.38	2.92	2.69	1.85	2.93	2.63	2.70	2.81	2.33	2.66	2.55	1.39	2.83	2.85
	61	62	63	64											
Precision	0.78	1.00	1.00	0.42											
Recall	0.35	0.55	0.30	0.50											
Payoff	0.90	0.92	0.85	0.91											
Entropy	2.01	2.27	1.61	2.80											

Table 3: The final shape categories extracted from asymmetric affinities. The number of outlier shapes is 80.

- [6] Anil K. Jain and Richard C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.
- [7] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks and deformations i: The components of two-dimensional shape and the reaction-diffusion space. *Int. J. Comput. Vision*, 15(3):189–224, 1995.
- [8] E. Klassen, A. Srivastava, W. Mio, and S. H. Joshi. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(3):372–383, 2004.
- [9] Y. J. Lee and K. Grauman. Foreground focus: Unsupervised learning from partially matching images. *International Journal of Computer Vision*, 2009.
- [10] H. Ling and D. Jacobs. Shape classification using the inner-distance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(2):286–299, 2007.
- [11] T. Liu and D. Geiger. Approximate tree matching and shape similarity. In *ICCV*, volume 1, pages 456–462, 1999.
- [12] D. Mumford. Mathematical theories of shape: Do they model perception? In B. C. Vemuri, editor, *roc. SPIE Vol. 1570, p. 2-10, Geometric Methods in Computer Vision*, pages 2–10, 1991.






























































	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
															
Precision	1.00	0.71	1.00	0.95	1.00	1.00	1.00	1.00	0.80	0.74	1.00	0.94	0.82	0.90	0.92
Recall	1.00	0.75	0.85	0.90	0.90	0.90	0.90	0.80	0.60	0.70	0.95	0.80	0.90	0.95	0.60
Payoff	0.96	0.95	0.93	0.94	0.95	0.94	0.96	0.96	0.91	0.95	0.95	0.96	0.96	0.95	0.94
Entropy	2.96	2.91	2.76	2.87	2.82	2.70	2.88	2.77	2.54	2.86	2.91	2.82	3.05	2.89	2.52
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
															
Precision	0.44	1.00	1.00	0.88	1.00	0.86	1.00	1.00	1.00	1.00	1.00	1.00	0.58	1.00	1.00
Recall	0.40	0.90	0.80	0.75	0.45	0.30	0.90	0.50	0.50	0.30	0.45	0.25	0.55	0.45	0.90
Payoff	0.94	0.96	0.93	0.95	0.93	0.86	0.94	0.93	0.93	0.87	0.87	0.80	0.93	0.93	0.95
Entropy	2.81	2.88	2.67	2.73	2.19	1.93	2.82	2.28	2.30	1.75	2.13	1.57	2.81	2.19	2.79
	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45
															
Precision	1.00	1.00	1.00	0.42	0.69	0.87	1.00	0.91	1.00	1.00	1.00	0.47	1.00	1.00	1.00
Recall	0.55	0.35	0.85	0.55	0.45	0.65	1.00	0.50	0.45	1.00	0.85	0.35	1.00	0.85	0.95
Payoff	0.92	0.91	0.94	0.95	0.93	0.93	0.96	0.93	0.92	0.97	0.94	0.90	0.95	0.96	0.97
Entropy	2.28	1.93	2.80	2.96	2.36	2.60	2.97	2.29	1.95	2.96	2.80	2.57	2.88	2.82	2.94
	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
															
Precision	1.00	1.00	0.95	0.67	1.00	0.50	1.00	1.00	0.91	0.74	1.00	1.00	0.52	1.00	1.00
Recall	0.90	0.65	1.00	0.70	0.75	0.25	0.95	0.50	1.00	0.70	0.20	0.90	0.60	0.50	0.40
Payoff	0.95	0.88	0.97	0.95	0.95	0.89	0.95	0.93	0.95	0.94	0.86	0.96	0.96	0.92	0.88
Entropy	2.82	2.21	3.00	2.85	2.70	2.06	2.86	2.28	2.86	2.70	1.39	2.83	3.07	2.29	1.99
	61														
															
Precision	0.33														
Recall	0.45														
Payoff	0.93														
Entropy	3.09														

Table 4: The final shape categories extracted by rendering the affinities symmetric. The number of outlier shapes is 43.

- [13] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(11):1105–1120, 1999.
- [14] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.
- [15] T. B. Sebastian, P. N. Klein, and B. B. Kimia. Recognition of shapes by editing shock graphs. In *ICCV*, volume 1, pages 755–762, 2001.
- [16] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000.
- [17] Alexander Strehl and Joydeep Ghosh. Cluster ensembles — a knowledge reuse framework for combining multiple partitions. *J. Mach. Learn. Res.*, 3:583–617, 2003.
- [18] A. Torsello and E. R. Hancock. Learning shape-classes using a mixture of tree-unions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(6):954–967, 2006.

- [19] A. Torsello, A. Robles-Kelly, and E. R. Hancock. Discovering shape classes using tree edit-distance and pairwise clustering. *International Journal of Computer Vision*, 72(3):259–285, 2007.
- [20] Andrea Torsello, Samuel Rota Bulo, and Marcello Pelillo. Grouping with asymmetric affinities: A game-theoretic perspective. In *CVPR*, volume 1, pages 292–299, 2006.
- [21] A. Tversky. Features of similarity. *Psychological Review*, 84:327–352, 1977.
- [22] J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA., 1995.
- [23] X. Yang, X. Bai, L. J. Latecki, and Z. Tu. Improving shape retrieval by learning graph transduction. In *ECCV*, pages 788–801, 2008.
- [24] S. C. Zhu and A. L. Yuille. Forms: A flexible object recognition and modeling system. *Int. J. Comput. Vision*, 20(3):187–212, 1996.