PROFI

Project number:	FP6-511572
Project acronym:	PROFI
Title:	Perceptually-relevant Retrieval Of Figurative Images

Deliverable No: D5.1: Algorithms for partial matching of regions

Short description:
Randomized algorithms are developed for finding good matches between parts of two shapes A and B. Shapes are modeled as sets of plane polygonal regions. The similarity (or dissimilarity) of matched parts is measured as the area of the symmetric difference.
Furthermore, like in work package 4 for the case of complete-complete matching, we apply the idea of matching based on image primitives to the problem of complete-partial matching. The approach is to decompose the shapes into geometric primitives, such as ellipses, triangles, rectangles and convex polygons, and to identify the relations between these objects. Then two shapes are matched based on the similarity between the objects and similarity of inter-object relationships. For the complete-partial matching the weights of unmatched objects of the partially matched figure are ignored.

Due month:	27
Delivery month:	27
Lead partner:	FUB
Partners contributed:	FUB
Classification:	Public

Project funded by the European Community under the "Information Society Technologies" Programme

1 Introduction

In this work package we develop algorithms for matching parts of two planar shapes. We assume that shapes are modeled by sets of plane simple polygons (regions).

As part of the work within this package we consider matching in a "classical" sense, that is finding a transformation that when applied to one shape minimizes a certain distance measure between two shapes. As possible classes of transformations we will consider translations, rigid motions (i.e., translations and rotations) and similarities (i.e., translations, rotations and scalings). The general situation in this setting is that we are given two objects A and B and a set T of allowable transformations and we want to transform B optimally so that the transformed image of B, or some part of it, is as close to A, or a part of A, as possible. Usually the quality of match is measured by some distance function or similarity measure $\delta(A, B)$ which assigns a number to any pair of objects A and B. We call the problem addressed in work package four of optimal matching of the complete shape B to the complete shape A a *complete-complete matching* (CCM). In this work package we consider the problem of *complete-partial matching* (CPM), i.e., matching B completely as good as possible to some part of A, and *partial-partial matching* (PPM), i.e. matching some part of B as good as possible to some part of A.

The area of overlap (or, equivalently, the area of the symmetric difference) of two planar regions is a natural measure of their similarity that is insensitive to noise. We analyze meaningful definitions of the partial matching problem in the context of trademark retrieval. Regarding complete-partial matching the natural variant of the distance measure when comparing the complete shape B to a part of A would be the area of $B \setminus A$. Minimizing this measure also means maximizing the area of overlap. However, especially for simple shapes this measure is not in all cases meaningful. Let for example shape A be a large circle and shape B be a small triangle that fits completely inside the circle. Putting B in A results in distance 0 although B would not be perceived as being similar to a part of A, so it is necessary to consider variants of the evaluation criteria and problem formulation. We discuss some possible variants of the distance measure that are appropriate for the problems of partial-complete and partial-partial matching. We also show that the probabilistic method described in deliverable report 4.1 is applicable for the partial matching problem as well.

We also apply the idea of matching based on decomposition into image primitives as described in work package 4 for the complete-complete matching problem to the problem of matching parts of the images. The main idea of this approach is to decompose an image into a set of primitives, such as ellipses, rectangles, triangles, or convex polygons. Parts of the image that do not belong to one of these simple classes are grouped together and classified as complex shapes. The identified objects get weights according to their size and "prägnanz", as understood in gestalt theory. Two images are compared based on the similarity between the primitives, the weights of the primitives and on their relative positions. The modification of the algorithm is fairly small for the problem of complete-partial matching: The weights of the objects in the partially matched image are ignored for the matching process and only influence the valuation of the matches. The extension of this algorithm to partial-partial matching is not as straightforward. We investigate possibilities to identify largest possible parts of the images that are similar under complete-complete matching.

2 Results

2.1 Probabilistic matching

For complete-complete matching of sets of regions we developed a probabilistic algorithm that finds a transformation minimizing the area of the symmetric difference of the two given sets if the allowed transformations are translations or rigid motions. The algorithm is described and analyzed in the deliverable report 4.1, here we shortly summarize the idea of the algorithm: We take a random sample of points of suitable size in each shape and record a "vote" for the transformation that maps the sample of shape A to the sample of shape B. The size of the sample depends on the class of the allowed transformations. If this experiment is repeated many times, we get a certain distribution of votes in transformation space. We showed that the density function of this distribution is proportional to the area of overlap of the two shapes induced by the corresponding transformation. The transformation with the largest number of votes in its neighborhood then approximately maximizes the area of overlap. For translations and rigid motions maximizing the area of overlap is equivalent to minimizing the area of the symmetric difference.

The area of the symmetric difference as a measure of dissimilarity of two shapes A and B indicates how much of the shape A is unmatched in B plus how much of B is unmatched in A. A natural extension of this idea to a complete-partial matching, if we want to match the complete shape A to some part of shape B, is to measure how much of the shape A is unmatched in B and to ignore the unmatched parts of B. That means for the complete-partial match we want to minimize the area of $A \setminus B$.

If the transformation does not allow scaling, which is the case with translations and rigid motions, minimizing the area of the difference $A \setminus B$ means to maximize the area of overlap of A and B. Therefore, the algorithm we described in deliverable 4.1 realizes complete-partial matching in a sense that it finds a transformation that approximately maximizes the area of overlap of two given sets of regions and, thus, minimizes the area of $A \setminus B$. For the cases of similarity maps and more general affine transformations analysis of the algorithm turns out to be quite complex and is part of our ongoing work. For the shapes that are similar in the sense that they contain parts with large possible areas of overlap, see for an example Figure 1, our algorithm finds the best transformation for the partial matching.

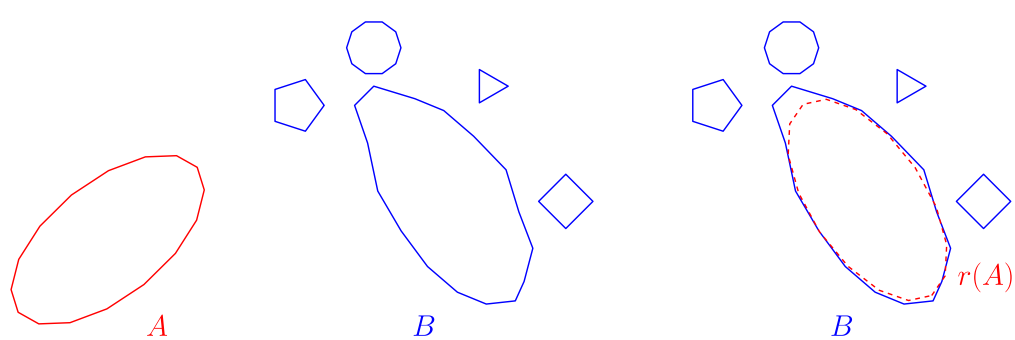

Figure 1: Partial matching of two shapes with large similar parts.

On the other hand, if the shape A is a small figure and the shape B has large parts where A can be completely fit into, see Figure 2 top row for an example, then any position of A

within the large part of B is equally good. Furthermore, even if B contains a region congruent to shape A (Figure 2, bottom row) the position of A on top of that region is as good as any position of A within the large part of B. The distance from A to B as measured by the area of $A \setminus B$ is zero in all described positions. Nevertheless, we would probably not consider A

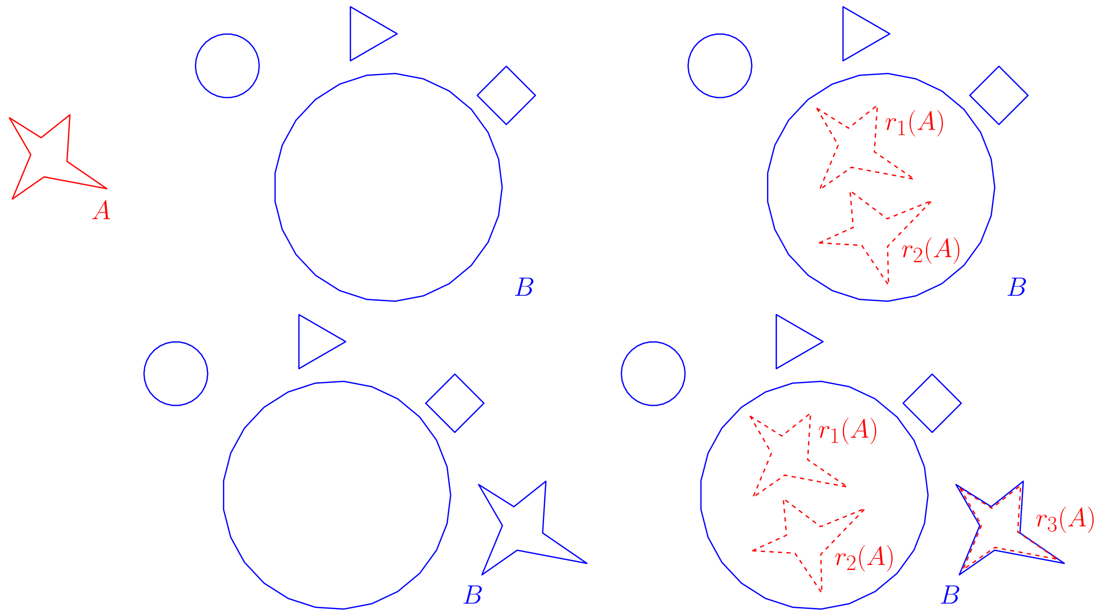

Figure 2: Partial matching of a small shape A to a shape B containing large parts.

as similar to the large part of B, but rather to the smaller nearly congruent part as match $r_3(A)$ indicates in the bottom row of Figure 2. This, with respect to the area of regions, natural notion of partial distance seems to be less suitable for perceptually relevant partial matching of shapes. It still remains interesting for other applications, for example, cutting stock problems.

For the perceptual similarity of the shapes defined by regions it is not only important how good one shape overlaps with some part of the other, but also how good the borders of the parts match. We plan to address this problem by considering a certain neighborhood of the shapes. The algorithm is modified as follows: The shapes are padded by an offset region (see Figure 3) and the random samples are generated in the interior regions of the shapes as well as in the offset region. We compute a transformation that maps one sample to the other, and record a positive vote for that transformation if both samples came from the interior of the shapes or both from the offset region, otherwise we record a negative vote. If a sample consists of more than one point, for example, two points in case of similarity transformations, we generate the first point randomly, and the following points either all from the interior of the shape if the first point in an internal one, or all from the offset region otherwise. Again, we take a transformation with highest vote count as the best matching transformation.

Let A' denote the shape A padded by its offset region and B' the padded shape B. The density of vote distribution in the transformation space of this modified algorithm is proportional to the area of overlap of the padded shapes reduced by twice the area of the symmetric difference restricted to the overlapping offset regions, formally $|A' \cap B'| - 2\,|(A \triangle B) \cap (A' \cap B')|$, see Figure 4.

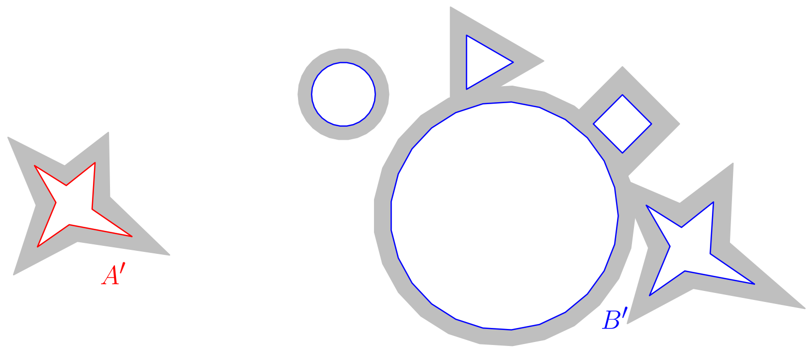

Figure 3: Shapes A and B padded by an offset region (shaded grey).

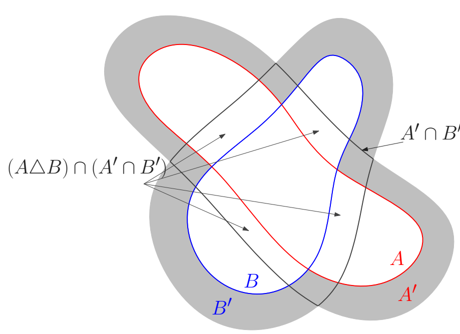

Figure 4: Similarity measure $|A' \cap B'| - 2\,|(A \triangle B) \cap (A' \cap B')|$ for partial matching.

Thus, we still want to maximize the area of overlap of A with some part of B, but also require the neighborhood of A to be similar to that of the matched part of B. This way in the previous example the positions of A within the large part of B would get less votes than the one where A coincides with the congruent part, see Figure 5. The detailed analysis of

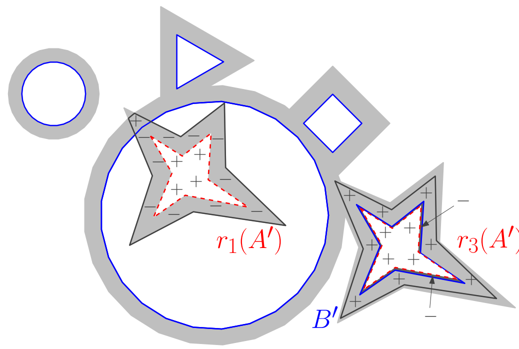

Figure 5: Partial matching with offset neighborhood.

this variant of the algorithm is still in progress. For the implementation of this algorithm we may consider taking the bounding boxes or minimum enclosing circle of the shapes for the extended neighborhood.

2.2 Matching based on Image Primitives

Definition of partial similarity

For the partial-partial-matching of figurative images based on sets of (polygonal) curves we developed a criterion for rating the similarity based on the similarity of the matched parts and on their size relative to the whole images. An analogous technique for the ppm problem based on regions is not applicable, because most of the visual information of a shape is coded in the boundary and not in the interior. When comparing a star and a circle, they do not appear to be similar just because parts of the star's interior do match well with parts of the circle's interior. However if the images are composed of several distinctive shapes, the similarity of these shapes can be rated and — together with the salience of these shapes — be a basis for the evaluation of partial similarity.

Matching Algorithm

In deliverable 4.1 we described an algorithm for the perceptually relevant comparison of figurative images (as given by sets of shapes). The idea in our approach is that an image is divided into a set P of (not necessarily spatially independent) parts — preferably simple and salient geometric figures. These parts are classified, weighted, and the set R of "relations" between the parts are identified. The relation are weighted as well based on the weights of the figures they connect. Comparing two images is accomplished by searching for subsets of the parts and their relations that match well.
The comparison of the parts is done independently, leaving aside their relative sizes and positions. It can be done using a similarity measure that works well for shapes whose parts lie close together whereas the resulting measure can handle arbitrary composed shapes.

For the comparison of two images I^1 and I^2 the relevance w_P of the figures and the relevance w_R of the relations is preset such that $w_P + w_R = 1$ — for images consisting only of one type of figures, e.g., only squares, the relations between these figures are of greater importance than for images consisting of totally different figures. The figures $p \in P$ and relations $r \in R$ get weights $w(p)$ and $w(r)$ such that for each image all weights sum up to 1, namely: $\sum_{p \in P} w(p) = w_P$ and $\sum_{r \in R} w(r) = w_R$.

For every pair $(p_i^1, p_k^2) \in P^1 \times P^2$ of figures and every pair $(r_{i,j}^1, r_{k,l}^2) \in R^1 \times R^2$ of relations, where $r_{i,j}$ denotes a relation between figures p_i and p_j, a value of similarity $s \in [0, 1]$ is computed, using simple measures of similarity.

Let $\mathcal{M}$ be the set of all one-to-one matchings between figures of image I^1 and image I^2. The value of similarity of the two images is then defined as the weighted sum of the similarities of the matched figures, plus the weighted sum of the similarities of the (implicitly) matched relations:

$$s(I^1, I^2) = \max_{M \in \mathcal{M}} \left\{ \sum_{(p^1,p^2) \in M} s(p^1, p^2) \cdot \frac{w(p^1) + w(p^2)}{2} + \sum_{\substack{(p_i^1,p_k^2) \in M \\ (p_j^1,p_l^2) \in M}} s(r_{i,j}^1, r_{k,l}^2) \cdot \frac{w(r_{i,j}^1) + w(r_{k,l}^2)}{2} \right\}$$

Finding images I^2 that contain a part that is similar to the query image I^1 is referred to as *partial-complete-matching*. This can easily be accomplished using our approach, by ignoring the weights of I^2. If *partial* refers to the matching between single figures, the similarity $s(p^1, p^2)$ of two figures p^1, p^2 has to be replaced by a partial-complete similarity measure $s_{\rightarrow}(p^1, p^2)$.

3 Deviations from plan

There have been no deviations from plan.

4 Project publications

Helmut Alt, Ludmila Scharf , Sven Scholz, Probabilistic Matching and Resemblance Evaluation of Shapes in Trademark Images, accepted for accepted for International Conference on Image and Video Retrieval, CIVR 2007.