

RuCCS TR-36

August, 1997

View-Based Object Recognition Using Saliency Maps

Ali Shokoufandeh


Center for Cognitive Science and
Department of Computer Science
Rutgers University
shokoufa@cs.rutgers.edu

Ivan Marsic

Department of Electrical and
Computer Engineering
Rutgers University
marsic@caip.rutgers.edu

Sven Dickinson

Center for Cognitive Science and
Department of Computer Science
Rutgers University
sven@cs.rutgers.edu



View-Based Object Recognition Using Saliency Maps*

Ali Shokoufandeh[§] Ivan Marsic[†] and Sven J. Dickinson[§]

[§]Department of Computer Science and
Center for Cognitive Science
Rutgers University
New Brunswick, NJ 08903
{shokoufa,sven}@cs.rutgers.edu


[†]Department of Electrical and
Computer Engineering
Rutgers University
New Brunswick, NJ 08903
marsic@caip.rutgers.edu

Abstract

We introduce a novel view-based object representation, called the *saliency map graph (SMG)*, which captures the salient regions of an object view at multiple scales using a wavelet transform. This compact representation is highly invariant to translation, rotation (image and depth), and scaling, and offers the locality of representation required for occluded object recognition. To compare two saliency map graphs, we introduce two graph similarity algorithms. The first computes the topological similarity between two SMG's, providing a coarse-level matching of two graphs. The second computes the geometrical similarity between two SMG's, providing a fine-level matching of two graphs. We test and compare these two algorithms on a large database of model object views.

Keywords: View-Based Object Recognition, Shape Representation and Recovery, Graph Matching.

*A condensed version of this paper will appear in the IEEE International Conference on Computer Vision, Bombay, January 1998, under the title, "View-Based Object Matching."



1 Introduction

The view-based approach to 3-D object recognition represents an object as a collection of 2-D views, sometimes called aspects or characteristic views [14]. The advantage of such an approach is that it avoids having to construct a 3-D model of an object as well as having to make 3-D inferences from 2-D features. Many approaches to view-based modeling represent each view as a collection of extracted features, such as extracted line segments, curves, corners, line groups, regions, or surfaces (Ikeuchi and Kanade [11], Burns and Kitchen [2], Ullman and Basri [30], Dickinson et al. [7], and Pope and Lowe [21]). The success of these view-based recognition systems depends on the extent to which they can extract their requisite features. With real images of real objects in unconstrained environments, the extraction of such features can be both time consuming and unreliable.

In contrast to the feature-based view-based recognition paradigm, a number of image-based view-based recognition systems have emerged. Beginning with the eigenface approach proposed by Turk and Pentland [29], these image-based approaches avoid extracting complex features from an image; instead, they retain the entire raw image as a single feature in a high-dimensional space. Turk and Pentland focused on the domain of faces and therefore did not require a large set of model views for each face. Nayar and Murase extended this work to general 3-D objects where a dense set of views was acquired for each object [19].

Although avoiding costly and often unreliable feature extraction, these image-based approaches pay the price of sensitivity to lighting conditions, image translation, image rotation, depth rotation, occlusion, and minor shape variation, all of which affect an image's pixel values and result in a change in the image's location in some high-dimensional space. Recent results have shown some progress towards solving these problems, e.g., the work of Belhumeur and Kriegman [1] (limited invariance to illumination changes) and the work of Leonardis and Bischoff [15] and Schmid and Mohr [24] (limited invariance to occlusion). Nevertheless, the lack of abstraction from raw image data to the model means that the model defines a very specific object instance.

The concept of computing coarse-to-fine image descriptions has much support in the

computer vision community; some examples include [3, 12, 16, 20, 27, 28]. In some cases, attention models have been developed that use a multiscale description to decide where in the image to apply some operation. Lindeberg has based this selection process on a quantitative analysis of gray-level blobs in scale space [16]. Jägersand [12] uses an information theoretic measure to compute “informativeness” of image regions at different scales, while others have defined some measure of “importance” and used it to drive an attention process [20, 27, 28]. Although suitable for locating objects in images for further processing, the above multiscale descriptions, often called saliency maps, lose the detailed shape information required for object recognition.

Some multiscale image descriptions have been used to locate a particular target object in the image. For example, Rao et al. use correlation to compare a multiscale saliency map of the target object with a multiscale saliency map of the image in order to fixate on the object [23]. Although these approaches are effective in finding a target in the image, they, like any template-based approach, do not scale to large object databases. Their bottom-up descriptions of the image are not only global, offering little means for segmenting an image into objects or parts, but offer little invariance to occlusion, object deformation, and other transformations.

An approach similar to the approach we will present is due to Crowley et al. [5, 4, 6]. From a Laplacian pyramid computed on an image, peaks and ridges at each scale are detected as local maxima. The peaks are then linked together to form a tree structure, from which a set of *peaks paths* are extracted, corresponding to the branches of the tree. During matching, correspondence between low-resolution peak paths in the model and the image are used to solve for the pose of the model with respect to the image. Given this initial pose, a greedy matching algorithm descends down the tree, pairing higher-resolution peak paths from the image and the model. Using a log likelihood similarity measure on peak paths, the best corresponding paths through the two trees is found. The similarity of the image and model trees is based on a very weak approximation of the trees’ topology and geometry, restricted, in fact, to a single path through the tree.

In this paper, we present a multiscale view-based representation of 3-D objects that, on

one hand, avoids the need for complex feature extraction, such as lines, curves, or regions, while on the other hand, provides the locality of representation necessary to support occluded object recognition as well as invariance to minor changes in both illumination and shape. In computing a representation for a 2-D image (whether model image or image to be recognized), a multiscale wavelet transform is applied to the image, resulting in a hierarchical saliency map of the image that offers advantages over a Laplacian pyramid. This saliency map is represented as a hierarchical graph structure, called the *saliency map graph*, that encodes both the topological and geometrical information found in the saliency map.

The similarity between a test image and a model image is defined as the similarity between their respective saliency map graphs. We address the problem of matching two saliency map graphs, leading to two matching algorithms. The first algorithm finds the best mapping between two saliency map graphs in terms of their topological structure, while the second algorithm factors in the geometry of the two graphs. In each case, we present an evaluation function that determines the overall quality of the match, i.e., the similarity of the two graphs. We demonstrate and evaluate our image representation and our two matching algorithms using the Columbia University COIL image database. In addition, we assess the viewpoint invariance of our representation and matching algorithms.

2 A Scale-Space Saliency Representation of an Image

To reduce the complexity in matching input image representations to model view representations, we seek a scale-space or coarse-to-fine representation of images that allows us to first match or index based on the coarse-level features in the image. Coarse-level correspondence can then be used to constrain a fine-level matching of the remaining features. Furthermore, we would like our image representation to be invariant to slight variations in the illumination falling on the object, image-plane rotation, translation, and scaling of the object, slight rotation in depth of the object, slight deformations of the shape of the object, e.g., stretching, bending, etc., and occlusion of the object.

Traditional view-based object representations that are image based, e.g., [29, 19, 1, 15,

24, 23], are neither coarse-to-fine nor invariant to the above transformations due to the global nature of their representations (although some offer limited invariance to particular transformations). However, the advantage of these approaches is that complex feature extraction, grouping, or abstraction is not required. Systems based on more invariant view-based image descriptions, e.g., [11, 2, 30, 7, 21], have relied on complex feature extraction (e.g., edges, lines, regions, etc.) which is not only unreliable but often requires domain-specific parameter tuning.

To address these shortcomings, we compute a scale-space representation of an image in which image objects (homogeneous regions) are located at the coarsest scale which captures their salient shape properties. Moreover, both the geometrical and topological relations between the regions will be explicitly encoded in the representation. Finally, computing these regions and relations requires the setting of very few parameters.

2.1 The Multiscale Wavelet Transform

The scale-space image representation that we have selected is based on a multiscale wavelet transform [25]. The advantage of the wavelet decomposition lies in its effective time (space)-frequency (scale) localization. Unlike other image transforms, e.g., [3, 4], which spread the information across their basis functions, the wavelet transform allows us to compute better localized object representations. In the output of the transform, as illustrated in Figure 1, the salient shape of small objects is best captured by small wavelets, while the converse is true for large objects. Searching from finer to coarser scales (right to left in Figure 1), we select the scale which captures the most efficient encoding of an object’s salient shape; above the chosen scale, extraneous information is encoded, while below the chosen scale, the object is overly “blurred.” The region defining the object at the chosen scale is called the *scale-space cell* (SSC) [18].

The dyadic wavelet transform of a function $f \in L^2(\mathbb{R})$ at the scale 2^j and at the position k is given by the inner products of the function with the family of wavelets

$$(Wf)(j, k) = \langle f, \psi_{j,k} \rangle = 2^{-j/2} \int_{-\infty}^{+\infty} f(x) \overline{\psi(2^{-j}x - k)} dx \quad (1)$$

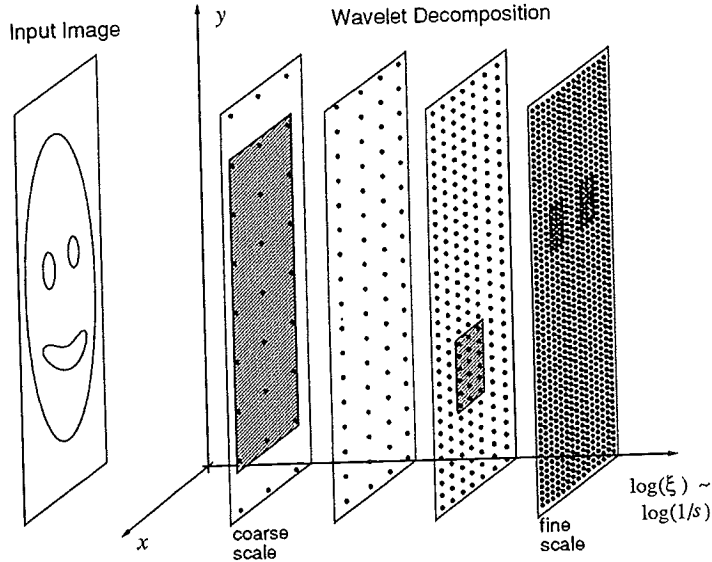


Figure 1: In the proposed multiscale description, objects are captured at the coarsest scale which captures their salient properties. Note that the frequency ξ and the scale s are inversely proportional.

where the overline denotes the complex conjugate. This inner product may also be viewed as a convolution product $(Wf)(j, k) = f * \overline{\psi_{j,k}}$ or as a filtering of the function f with a band-pass filter whose impulse response is $\overline{\psi_{j,k}}$.

Detecting the scale-space cells requires analysis of the wavelet transform response at each scale. The scale-space cell of an image object is located at the scale which is approximately one octave below the scale at which the object's response becomes indistinguishable from other image objects of the same size. At this scale, the object's response resembles the wavelet basis function impulse response. The following subsections will explore scale-space cell detection in greater detail.

2.2 Scale-Space Cells in One Dimension

To illustrate the detection of scale-space cells, consider the one-dimensional signal in Figure 2, using the wavelets described in [17]. Any object behaves like a point at all scales coarser than its characteristic scale I . This means that the wavelet transform of a signal at any

scale $j > I$ is the same no matter what the object's shape is; the transform is completely determined by the width of the objects and their amplitude. We can ignore the information at the scales $j > I$ and still be able to reconstruct the original signal almost perfectly, as illustrated in the right column in Figure 2.

The reconstructed signal $f_6^r(x)$ is obtained by replacing the responses $(W\cdot)(6, x)$ of objects A and B with the analyzing wavelet $\psi_{6,k}$ of the same amplitude. Similarly, $f_{6,5}^r(x)$, $f_{6,5,4}^r(x)$, and $f_{6,5,4,3}^r(x)$ are obtained by successively replacing the corresponding wavelet responses at the scales (6,5), (6,5,4), and (6,5,4,3), respectively, with the analyzing wavelets. The error in reconstruction is very small if the replacement takes place at scales greater than the characteristic scale; the error increases significantly as the replacement takes place at the finer scales. In the general case, interactions between the neighboring objects will distort the wavelet transform response. However, even for complex signals, each object will eventually yield the (approximate) impulse response at the appropriate scale determined by the size of the object.

In order to find the characteristic scale of an image object in the 1-D case, one can measure the correlation between the wavelet transform of the object and the basis function at any given scale. At each location (x, y) , one can then select the finest scale at which the correlation exceeds some threshold. We will now proceed to examine the detection of a scale-space cell in two dimensions.

2.3 Saliency Detection Algorithm in Two Dimensions

In the two-dimensional case, the characteristic scale $I(\Theta)$ may be different for any particular orientation Θ of a 1-D cross-section through an object. Any object other than a circular one (disc) will become a point at different scales in different directions, e.g., an elongated object, occluded object, etc. In this case, the object will extend over several SSC's at any scale $j < \max_{\Theta}\{I(\Theta)\}$. Therefore, in the 2-D case, we apply the 1-D procedure in a number of directions and search for clusters of 1-D centers, as shown in Figure 3(a). Our entire procedure for detecting the SSC's in an image therefore consists of the following four steps [18]:

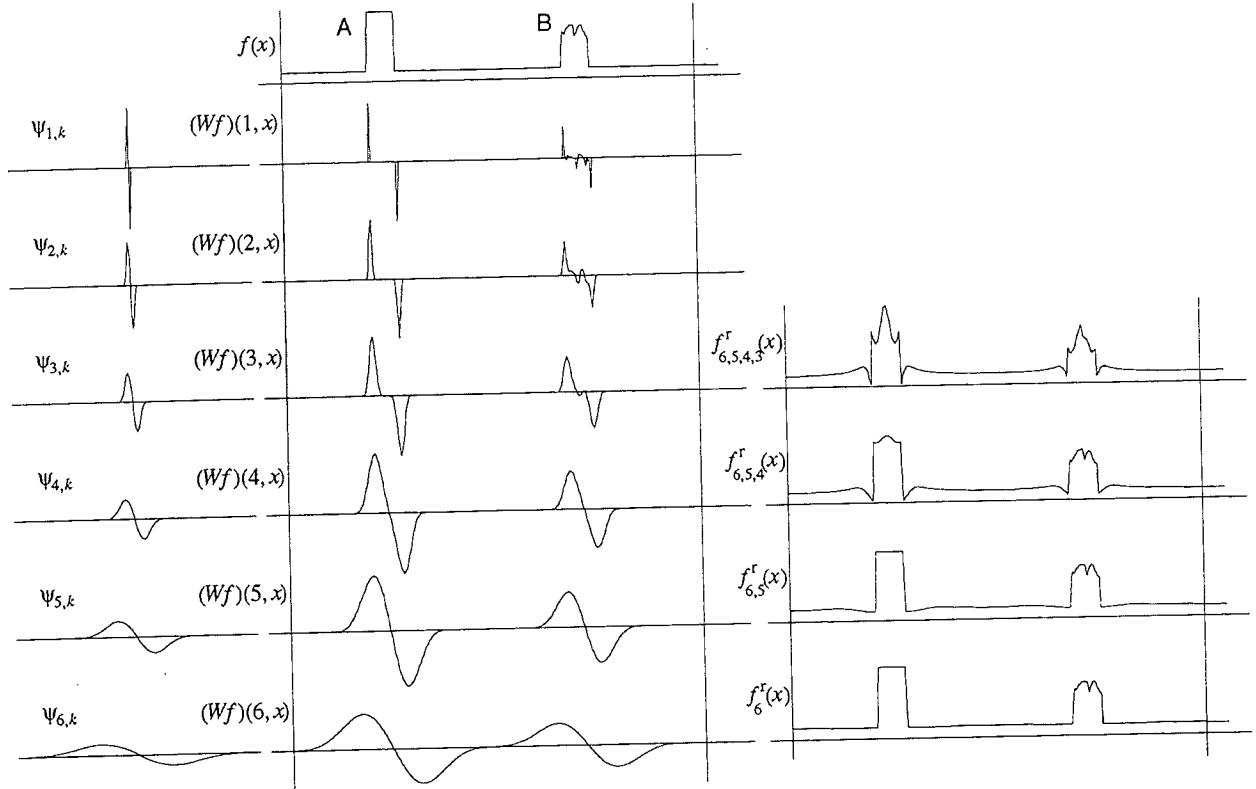


Figure 2: An illustration of the object's characteristic scale. The left column shows the analyzing wavelets $\psi_{i,k}$ at the corresponding scales. The middle column shows two examples (A and B) of a function $f(x)$ (top) and their wavelet transforms (below). Beginning with the characteristic scale $I = 4$, the $(W\cdot)(I + j, x)$ takes the shape of $\psi_{I+j,k}$, $j = 1, 2, \dots$. The right column shows the reconstruction of the original function by replacing the $(W\cdot)(i, x)$ with the $\psi_{i,k}$.

Step 1—Wavelet Transform: Compute the wavelet pyramid of an image with ℓ dyadic scales using oriented quadrature bandpass filters tuned to 16 different orientations, i.e. $\Theta = 0^\circ, 22.5^\circ, 45^\circ, \dots, 337.5^\circ$. See [26] for a detailed derivation and description of computing the wavelet pyramid using steerable basis filters.

Step 2—Local Energies: Compute the oriented local energies using the equation:

$$E(\Theta, s, x, y) = [G^\Theta(s, x, y)]^2 + [H^\Theta(s, x, y)]^2 \quad (2)$$

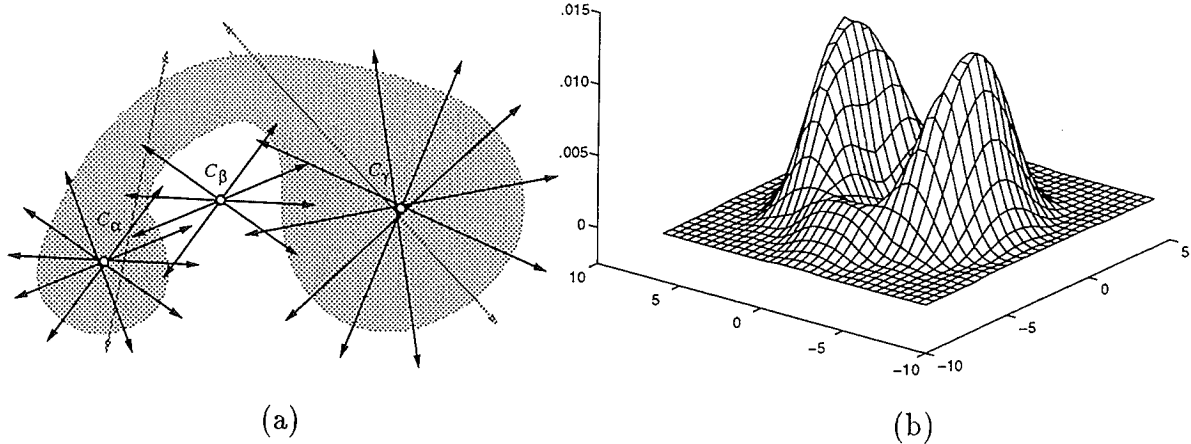


Figure 3: (a) The clusters formed by the centers of the 1-D SSCs associated with the cross-sections through an object. (b) One of the filter kernels $\vartheta(\Theta, x, y)$ used in computing the saliency of the SSC's (Eq. (3)). The kernel is obtained by computing the oriented energy at $\Theta = 0^\circ$ for a disc.

where $G^\Theta(s, x, y)$ and $H^\Theta(s, x, y)$ are the outputs of a quadrature pair of analyzing wavelet filters at the scale-space coordinate (s, x, y) , oriented at the angle Θ . For each image point, 16 different oriented local energies are computed.

Step 3—Saliency Maps: Compute ℓ saliency maps. The saliency of each particular SSC is computed using the convolution:

$$\text{saliency SSC}(s, x, y) = \sum_{\Theta} [E(\Theta, s, x, y) * \vartheta(\Theta, x, y)] \quad (3)$$

where $\vartheta(\Theta, x, y)$ is the filter kernel obtained by computing the sum of the squared impulse responses of the two analyzing wavelet filters $G^\Theta(s, x, y)$ and $H^\Theta(s, x, y)$, as shown in Figure 3(b). As discussed above, circular shape has the highest saliency as measured by this scheme.

Step 4—Peaks in Saliency Maps: Moving from finer to coarser scales at every location, we select the first saliency map for the which a peak (local maximum) at that location exceeds a given threshold. By using a series of oriented 1-D filters to detect

the characteristic scale, we can detect objects that are not perfectly circular in shape. For example, if a non-circular shape's variation in diameter does not reach neighboring scales above or below the current scale, then a circularly-symmetric filter, such as that used by Crowley [5, 4, 6], will give a weak response for the shape. In our approach, however, the 1-D filters are slightly adjusted in width (bounded by neighboring scales). The result is a cluster of oriented peaks from which we compute the 2-D shape's location as the centroid of these peaks. The saliency of the 2-D shape is computed as the sum of the oriented saliencies of the oriented peaks near this centroid. Finally, we apply a non-maximum suppression process to eliminate closely overlapping salient SSC's at each scale.

The contents of each scale-space cell (SSC) is a 2-D matrix of wavelet coefficients. The size of this matrix is invariant to both the scale at which the SSC is detected and the complexity of the shape contained in the SSC's corresponding image region. In the current implementation, only the scale, position, and saliency of a SSC is exploited during the matching of two saliency maps. However, one could include the actual content of the SSC as specified by the wavelet coefficient matrix. The fact that all the SSC matrices are self-similar and small (in our case, 16×16) means that efficient comparisons can be made between SSC's at different scales.

Figure 4 illustrates how a complex object's saliency map (a) is largely invariant to scaling (b), translation (c), image plane rotation (d), and limited rotation in depth (e), where the illuminated left side of the face exhibits little change in its saliency map. Circles in the image correspond to scale-space cells, while their intensity is proportional to their saliency. Note that the size of the circle appears to be slightly larger than its corresponding image feature. This is due to the fact that the size of the circle is determined by the largest extent of the filter shown in Figure 3(b), i.e., where the response approaches zero. For different images taken under different conditions, there cannot be true invariance in the sense that salient regions are identical in different views. However, approximate invariance suffices for the recognition scheme that we propose below.

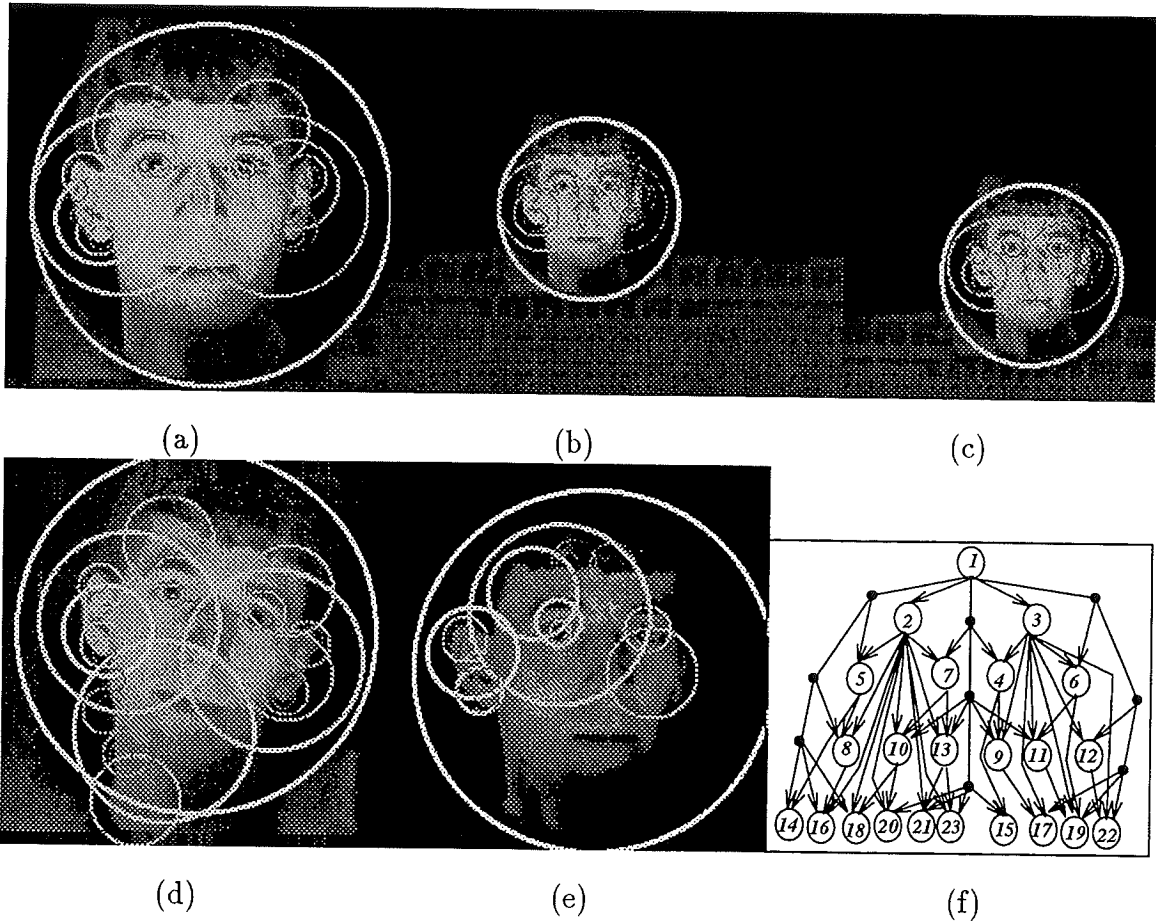


Figure 4: Extracting the most salient SSC's in an image: (a) original image and its saliency map; (b) scale invariance; (c) translation invariance; (d) image rotation invariance; (e) invariance to rotation in depth (illuminated left side of face exhibits little change in its saliency map); and (f) the saliency map graph (SMG) of the original image in (a).

2.4 Limitations of the Representation

Under normal circumstances, an object (or one of its component features) should produce a peak at its corresponding location in the saliency map. However, there are several exceptions which do not properly fit within the SSC framework. The evolution of the saliency map at a single scale as the object becomes degenerate, in the sense of the SSC framework, is shown in Figure 5. The exceptions arise due to wavelet transform scale sub-sampling, crowded objects, and elongated objects. As expected, combinations of these will create even more difficulties.




Figure 5: The evolution of the saliency map at a single scale for the exception cases. The left side shows the original images and the right side shows the corresponding saliency maps at one scale. (a) As the object size increases, a peak in the saliency map corresponding to the object turns into a “crater.” (b) As the objects approach closer to one another, the “anti-object” between them becomes the most salient. (c) Instead of peak(s) in the saliency map, elongated objects produce “mountain ridges.”

As shown in Figure 5(a), the saliency peak detector will find one salient region for the object at the left side, whereas it will find several salient regions for larger objects on the right side. This effect will occur when the size of an object is intermediate between two (octave) scales, and can be minimized by increasing the number of scales. In Figure 5(b), the detector will find four salient regions for the distant objects at left, but will find only one region in the center of the four close objects at right. This phenomenon occurs when a set of objects enclose a compact background region; the detector cannot separate figure from ground. One could argue that such a regular pattern represents a form of texture that should be treated as a single object. Again, if there were more scales, this composite object would be detected at a coarser scale. Finally, in Figure 5(c), the detector will find several salient regions positioned along the elongated object. In this case, a salient region grouper could search for a string of co-curvilinear SSC’s at a given scale and group them into a composite structure which could easily be accommodated by our graph representation and matching

Given this formulation of the mapping, f , we define the error of f to be:

$$\mathcal{E}(f) = \varepsilon \sum_{u \in V_1} \sum_{v \in V_2} M_{u,v} \omega(u, v) |s(u) - s(v)| + (1 - \varepsilon) \left(\sum_{u \in V_1} (1 - \sum_{j \in V_2} M_{u,j}) s(u) + \sum_{v \in V_2} (1 - \sum_{i \in V_1} M_{i,v}) s(v) \right) \quad (4)$$

where $\varepsilon = |\mathbf{1}^t M(f) \mathbf{1}| / |V_1|$ represents the ratio of the number of matched vertices to the number of vertices ($|V_1|$) in the model SMG (with $\mathbf{1}$ the identity vector of the appropriate dimension) and $s(v)$ denotes the strength of region v in its saliency map. For the SMG topological similarity algorithm, defined in Section 3.3, $\omega(u, v)$ is always one, while for the SMG geometrical similarity algorithm, defined in Section 3.4, $\omega(u, v)$ represents the Euclidean distance between the centers of the regions, u and v . Clearly, in the case of perfect similarity, $\mathcal{E}(f) = 0$, while $\mathcal{E}(f)$ will be $\sum_{u \in V_1} s(u) + \sum_{v \in V_2} s(v)$ if there is no match ($\mathbf{1}^t |M(f)| \mathbf{1} = 0$).

3.3 A Matching Algorithm Based on Topological Similarity

In this section, we describe an algorithm which finds an approximate solution to the SMG similarity problem. The focus of the algorithm is to find a minimum weight matching between vertices of G_1 and G_2 which lie in the same level. Our algorithm starts with the vertices at level 1. Let A_1 and B_1 be the set of vertices at level 1 in G_1 and G_2 , respectively. We construct a complete weighted bipartite graph $G(A_1, B_1, E)$ with a weight function defined for edge (u, v) ($u \in A_1$ and $v \in B_1$) as $w(u, v) = |s(v) - s(u)|$.² Next, we find a maximum cardinality, minimum weight matching M_1 in G using [9]. All the matched vertices are mapped to each other; that is, we define $f(x) = y$ if (x, y) is a matching edge in M_1 .

The remainder of the algorithm proceeds in phases as follows, as shown in Figure 6. In phase i , the algorithm considers the vertices of level i . Let A_i and B_i be the set of vertices of level i in G_1 and G_2 , respectively. Construct a weighted bipartite graph $G(A_i, B_i, E)$ as follows: (v, u) is an edge of G if either of the following is true: (1) Both u and v do not have any parent in G_1 and G_2 , respectively, or (2) They have at least one matched parent

² $G(A, B, E)$ is a weighted bipartite graph with weight matrix $W = [w_{ij}]$ of size $|A| \times |B|$ if, for all edges of the form $(i, j) \in E$, $i \in A$, $j \in B$, and (i, j) has an associated weight $= w_{i,j}$.

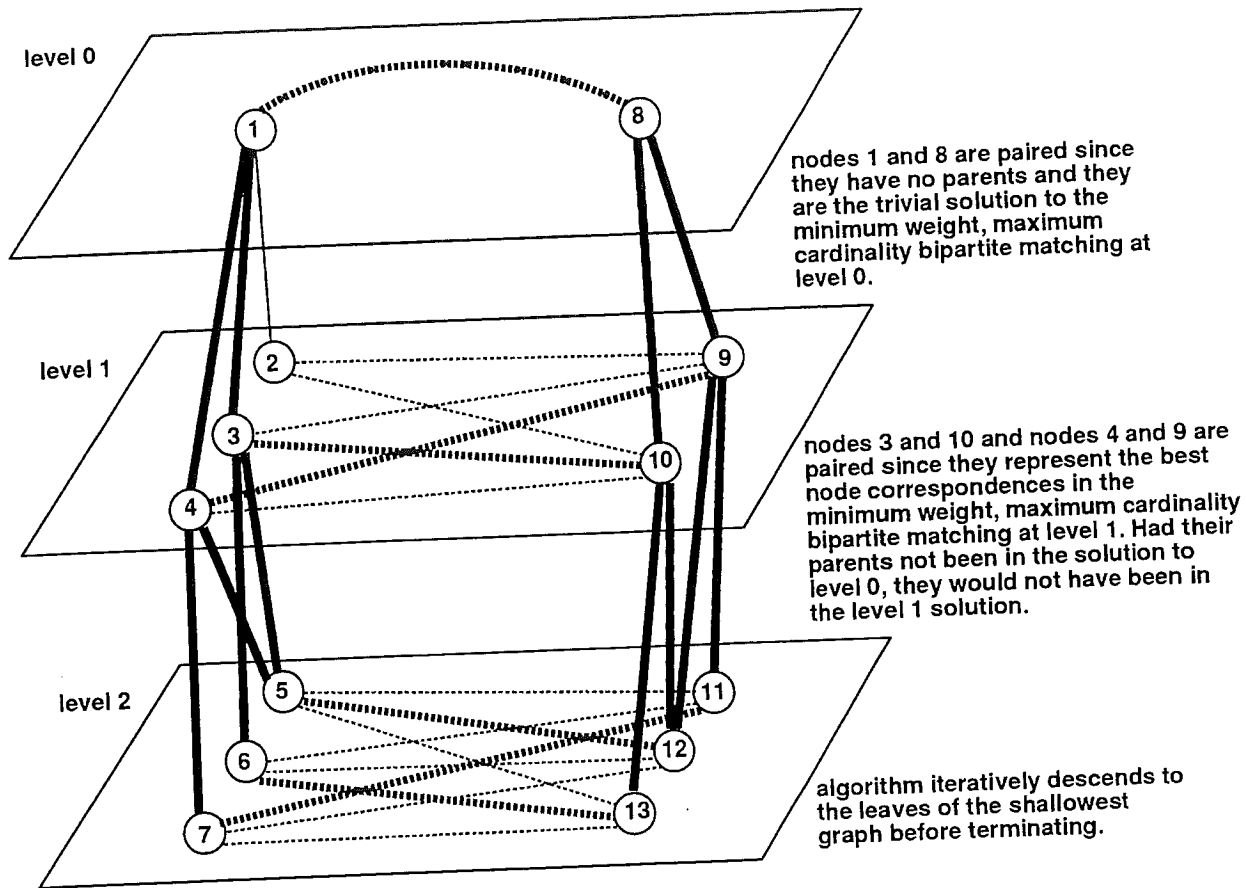


Figure 6: Illustration of the SMGBM Algorithm (see text for explanation).

of depth less than i ; that is, there is a parent p_u of u and p_v of v such that $(p_u, p_v) \in M_j$ for some $j < i$. We define the weight of the edge (u, v) to be $|s(u) - s(v)|$. The algorithm finds a maximum cardinality, minimum weight matching in G and proceeds to the next phase.

The above algorithm terminates after ℓ phases, where ℓ is the minimum number of scales in the saliency maps (or SMG's) of two graphs. The partial mapping M of SMG's can be simply computed as the union of all M_i 's for $i = 1, \dots, \ell$. Finally, using the error measure defined above, we compute the error of the partial mapping M . Each phase of the algorithm requires simple operations with the time to complete each phase being dominated by the time to compute a minimum weight matching in a bipartite graph. The time complexity for finding such a matching in a weighted bipartite graph with n vertices is $O(n^2 \sqrt{n \log \log n})$ time, using the scaling algorithm of Gabow, Gomans and Williamson [10]. The entire procedure,

as currently formulated, requires $O(\ell n^2 \sqrt{n \log \log n})$ steps.

3.4 A Matching Algorithm Based on Geometric Similarity

The SMGBM similarity measure captured the structural similarity between two SMG's in terms of branching factor and node saliency similarity; no geometric information encoded in the SMG was exploited. In this section, we describe a second similarity measure, called SMG Similarity using an Affine Transformation (SMGAT), that includes the geometric properties (e.g., relative position and orientation) of the saliency regions.

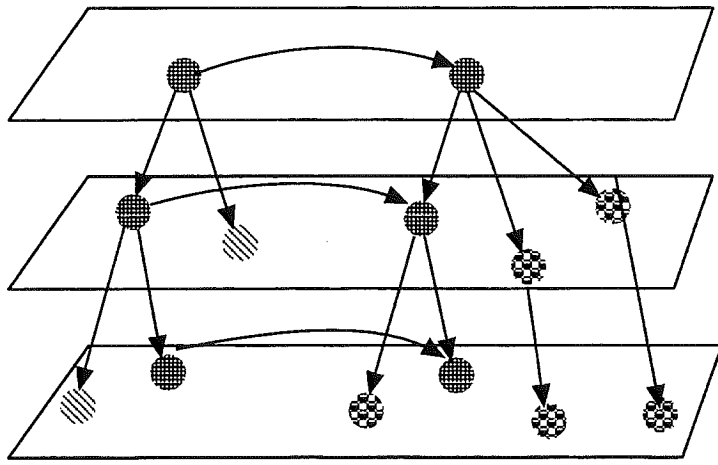
Given $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, we first assume, without loss of generality, that $|V_1| \leq |V_2|$. First, as shown in Figure 7, the algorithm will hypothesize a correspondence between three regions of G_1 , say (r_1, r_2, r_3) , and three regions (r'_1, r'_2, r'_3) of G_2 . The mapping $\{(r_1 \rightarrow r'_1), (r_2 \rightarrow r'_2), (r_3 \rightarrow r'_3)\}$ will be considered as a basis for alignment if the following conditions are satisfied:

- r_i and r'_i have the same level in the SMG's, for all $i \in \{1, \dots, \ell\}$.
- $(r_i, r_j) \in E_1$ if and only if $(r'_i, r'_j) \in E_2$, for all $i, j \in \{1, \dots, \ell\}$, which implies that selected regions should have the same adjacency structure in their respective SMG's.


Once regions (r_1, r_2, r_3) and (r'_1, r'_2, r'_3) have been selected, we solve for the affine transformation (A, b) , that aligns the corresponding region triples by solving the following system of linear inequalities:

$$\begin{bmatrix} x_{r_1} & y_{r_1} & 1 & 0 & 0 & 0 \\ x_{r_2} & y_{r_2} & 1 & 0 & 0 & 0 \\ x_{r_3} & y_{r_3} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_{r_1} & y_{r_1} & 1 \\ 0 & 0 & 0 & x_{r_2} & y_{r_2} & 1 \\ 0 & 0 & 0 & x_{r_3} & y_{r_3} & 1 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ b_1 \\ a_{21} \\ a_{22} \\ b_2 \end{bmatrix} = \begin{bmatrix} x_{r'_1} \\ x_{r'_2} \\ x_{r'_3} \\ y_{r'_1} \\ y_{r'_2} \\ y_{r'_3} \end{bmatrix}. \quad (5)$$

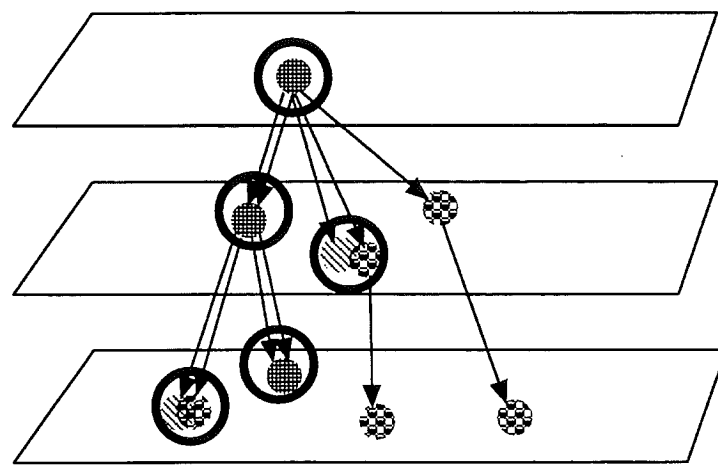
The affine transformation (A, b) will be applied to all regions in G_1 to form a new graph G' . Next, a procedure similar to the minimum weight matching, used in the SMGBM is applied to



Choose region triple correspondence and solve for affine transformation T that aligns the triples.



Apply affine transform T to one graph, aligning it with the other.



Find minimum weight mapping in bipartite graph at each level based on Euclidean distance.

Figure 7: Illustration of the SMGAT Algorithm (see text for explanation)

the regions in graphs G' and G_2 . Instead of matching regions which have maximum similarity in terms of saliency, we match regions which have minimum Euclidean distance from each other. Given two regions u and v , the distance between them can be defined as the L_2 norm of the distance between their centers, denoted by $d(u, v) = \sqrt{(x_u - x_v)^2 + (y_u - y_v)^2}$. In a series of steps, SMGAT constructs weighted bipartite graphs $\mathcal{G}_i = (R_i, R'_i, E_i)$ for each level i of the two SMG's, where R_i and R'_i represent the set of vertices of G' and G_2 at the i -th level, respectively. The constraints for having an edge in E_i is the same as SMGBM: (u, v) is an edge in \mathcal{G}_i if either of the followings holds:

- Both u and v do not have any parents in G' and G_2 , respectively.
- They have at least one matched parent of depth less than i .

The corresponding edge will have weight equal to $w(u, v) = d(u, v)$. A maximum cardinality, minimum weight bipartite matching M_i will be found for each level \mathcal{G}_i , and the partial mapping $f_{(A,b)}$ for the affine transformation (A, b) will be formed as the union of all M_i 's. Finally, the error of this partial mapping $\mathcal{E}(f_{(A,b)})$ will be computed as the sum over each E_i of the Euclidean distance separating E_i 's nodes weighted by the nodes' difference in saliency. Once the total error is computed, the algorithm proceeds to the next valid pair of region triples. Among all valid affine transformations, SMGAT chooses that one which minimizes the error of the partial mapping.

In terms of algorithmic complexity, solving for the affine transformation (eq. 5) takes only constant time, while applying the affine transformation to G_1 to form G' is $O(\max(|V_1|, |E_1|))$. The execution time for each hypothesized pair of region triples is dominated by the complexity of establishing the bipartite matching between G_2 and G' , which is $O(\ell n^2 \sqrt{n \log \log n})$, for SMG's with n vertices and ℓ scales. Although in the worst case, i.e., when both saliency map graphs have only one level, there are $O(n^6)$ pairs of triples. However, in practice, the vertices of an SMG are more uniformly distributed among the levels of the graph, greatly reducing the number of possible correspondences of base triples.

Although not yet implemented, we can reduce the complexity of the bipartite matching step by exploiting the fact that the edge weights of the bipartite graph represent the Eu-