



Adaptive Hough transform for the detection of natural shapes under weak affine transformations

Olivier Ecabert^{a,b,*}, Jean-Philippe Thiran^c

^a Philips Research, Weisshausstrasse 2, 52066 Aachen, Germany

^b Darmstadt University of Technology, Merckstr. 25, 64283 Darmstadt, Germany

^c Signal Processing Institute (STIITS), Swiss Federal Institute of Technology (EPFL), 1015 Lausanne, Switzerland

Received 26 February 2003; received in revised form 8 April 2004

Available online 17 June 2004

Abstract

This paper introduces a two-steps adaptive generalized Hough transform (GHT) for the detection of non-analytic objects undergoing weak affine transformations in images. The first step of our algorithm coarsely locates the region of interest with a GHT for similitudes. The returned detection is then used by an adaptive GHT for affine transformations. The adaptive strategy makes the computation more amenable and ensures high accuracy, while keeping the size of the accumulator array small. To account for the deformable nature of natural objects, local shape variability is incorporated into the algorithm in both the detection and reconstruction steps. Finally, experiments are performed on real medical data showing that both accuracy and reasonable computation times can be reached.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Adaptive; Generalized Hough transform; Affine transformation; Object detection; Local shape variability

1. Introduction

Detecting objects in images is a very active field of computer vision. Applications are numerous such as pattern recognition in industrial processes, detection of anatomical structures in medical images or object localization in aerial or satellite data. Most of the usual segmentation algorithms rely on very *local* image information (edges, pixel

gray levels) and use to fail if the initialization is performed too far away from the expected solution. For instance, the widely used active contours (Kass et al., 1998) are very myopic since the curve propagates according to a differential equation, that is, the particles of the snake are moved under the influence of a very small neighborhood of image pixels. If an application requires full automatic object extraction, *global* information about the structure of interest (e.g. shape) has to be encoded in the segmentation algorithm. However, most of the methods relying only on global constraints do not reflect shape variations at small scales. Approaches combining both local and global types of

* Corresponding author. Address: Philips Research, Weisshausstrasse 2, 52066 Aachen, Germany.

E-mail address: olivier.ecabert@a3.epfl.ch (O. Ecabert).

information are usually a kind of trade-off between global shape alignment and fidelity to local image features (e.g. Chakraborty et al., 1996; Cootes et al., 1994; Leventon et al., 2000).

In order to address this problem, we propose a two-steps approach, in which the structure of interest is first globally and rather coarsely detected using a Hough transform. This estimate can be then refined, in a second step, with any well-known local algorithm. For instance, the use of a geodesic active contour model (Ecabert and Thiran, 2002) makes the segmentation free of global constraints, providing good convergence to the object boundary.

1.1. Hough transform

Initially designed for analytic curves (Hough, 1962), the Hough transform (HT) was then generalized to non-analytic shapes by Ballard, 1981. This method is able to detect any arbitrary shape undergoing a geometric transformation in an image. Moreover, it has shown to be robust and can even be used successfully for the detection of overlapping or semi-occluded objects in noisy images. However, increasing the number, the range or the accuracy of the parameters of the geometric transformation may lead to high computation efforts, which are practically not amenable.

In the context of straight lines detection, Illingworth and Kittler (1987) proposed to implement the HT efficiently using an adaptive accumulator array and a coarse-to-fine strategy, so that only the areas of interest are investigated in greater details. Obvious advantages of this approach are that the process can go on until a given accuracy is reached without increasing the size of the array.

The randomized HT (RHT) offers a different approach to achieve increased efficiency in the detection of analytic curves (Xu and Oja, 1993), whereby a set of n pixels is randomly selected from the edge image, determining the n parameters of the curve of interest. This yields the coordinates of a single point in the parameter space and therefore only one cell of the accumulator array needs to be incremented. Besides, the dynamic implementa-

tion of the accumulator array along with the “exact” computation of the parameters provide with infinite range and high resolution of the parameters.

Straight lines and analytic curves have a limited range of applications, especially for real life scenes or medical image analysis. The shapes to detect are usually provided as a pixel map of typical appearances of the object of interest. Based on the pioneer work of Ballard (1981), many algorithms have been proposed to detect more general shape undergoing geometric (e.g. similarity or affine) transformations. The use of the gradient directional information acts favorably in two points (Montiel et al., 2001): it increases the detection speed and improves the parameter accuracy by reducing the wrong evidences. However, obvious problems still reside if the number of degrees of freedom of the geometric transformation is high. Indeed, the six parameters of the affine transformation would produce a six-dimensional accumulator array. Moreover, the entire process has to be repeated for every combination of the discrete parameter values. If wide range and high accuracy are required for the parameters, both computational time and storage space become very large. These methods are also referred to as brute force voting, or 1-to- n mapping, approaches. Indeed, one point of the edge image produces the incrementation of n cells of the accumulator array.

They are to compare to the so-called n -to-1 mapping methods (e.g. Aguado, 2002; Kimura and Watanabe, 2002; Ser and Siu, 1995; Fung et al., 1996), which from n image points produce only one incrementation in the array. They are usually fully invariant respective to the underlying geometric transformation and the parameters can be computed at once from image points, reducing therefore the number of superfluous operations and the computation time. However, according to our experience, the techniques used to extract geometric invariant points from the image are very sensitive to noise and object overlapping. Moreover, the results are further degraded by the fact that two natural shapes are not supposed to be exactly in geometric correspondence.

1.2. Proposed approach

This paper is concerned with the description of a generalized HT for the extraction of natural shapes from images. Unlike the original generalized HT (GHT) mainly designed for similarity transformations introduced by Ballard (1981), we propose a two-steps approach able to deal with affine transformations. Indeed, the wide range of variability that natural shapes can cover is better described by a geometric transformation having six parameters than four.

However, the six parameters of the affine transformation would produce a six-dimensional array, leading to unrealistic computational time and storage requirement, if high parameter accuracy is expected. For that purpose, we propose a two-steps adaptive approach. In the first step, a standard GHT for similitudes with poor discretization of the accumulator array axis is performed on the image. In the second step, the resulting coarse detection is then applied as initialization to an adaptive GHT for affine transformations. The use of an adaptive scheme makes it possible to achieve high accuracy without increasing the size of the accumulator array. This method will be referred to as HT for shapes under weak affine transformations.

Another major contribution of this paper is the incorporation of local shape variability at three stages of the detection. First, natural shapes are considered by encoding in the R-table not only one reference shape, but several instances of the object of interest. Second, an uncertainty region is associated to each entry index of the R-table, accounting for different sources of error influencing the computation of the shape parameters. This uncertainty region can be directly derived from the generalized R-table introduced previously. Finally, a priori shape information is used intelligently along with the gradient image to reconstruct the object from the peak of the accumulator array.

This paper is organized as follows. In Section 2 the two-steps HT for shapes under weak affine transformations is introduced. Section 3 describes the different techniques used to account for natural shape variability in the detection process. Finally,

experimental results on magnetic resonance images of the brain are shown in Section 4.

2. Hough transform for shapes under weak affine transformations

n -to-1 mapping methods have the advantage of being fully affine invariant, but are unfortunately not robust for real images and for the recognition of shapes that are not exactly affine transformed. Approaches using brute force voting like the GHT can overcome this problem by increasing the space complexity. Indeed, for each combination of the transformation parameters and for each edge pixel, all the corresponding cells of the accumulator array need to be incremented. Nevertheless, the gradient direction can still be used to reduce the number of false evidences, and errors due to discrete computation can be taken into account as proposed by Aguado et al. (2000).

However, if a shape under affine transformation is to be detected, the GHT has to be repeated M^4 times, where M is the number of discrete values for each parameter of the linear transformation. Direct implementation cannot be carried out if time performances are required. For that purpose, we propose a two-steps implementation, where in the first step, a GHT for similarity transformations with poor axis discretization is performed on the image, leading to a coarse detection. Complexity increases thus as M^2 if both independent parameters (i.e. isotropic scale and orientation) are discretized with M values. The obtained result is then used, in the second step, as initialization to an adaptive GHT for affine transformations, where each axis of the parameter space is discretized with only three values. The range of the axis values is re-estimated iteratively during the coarse-to-fine search strategy.

2.1. Generalized Hough transform for similarity transformations

To introduce the GHT, let W be a shape template and \mathbf{x} be a point of the gradient image. Moreover, let us define a geometric transformation of the shape template by

$$\mathbf{A}W + \mathbf{t} = \begin{bmatrix} a_A & b_A \\ c_A & d_A \end{bmatrix} \cdot \begin{bmatrix} W_x \\ W_y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix}$$

$$(a_A d_A - b_A c_A \neq 0),$$

with \mathbf{A} and \mathbf{t} corresponding respectively to a linear transformation and to a translation vector. In the affine case, \mathbf{A} can be interpreted as the product of elementary matrices ($\mathbf{A} = \mathbf{A}_i \cdot \mathbf{A}_j, \dots$), corresponding to elementary geometric operations like rotations \mathbf{A}_r , non-uniform scalings \mathbf{A}_s and shearings \mathbf{A}_{sh} with

$$\mathbf{A}_r = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}, \quad \mathbf{A}_s = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix},$$

$$\mathbf{A}_{sh} = \begin{bmatrix} 1 & sh_x \\ sh_y & 1 \end{bmatrix}.$$

The similarity transformation is a particular case of the affine transformation with $sh_x = sh_y = 0$ and $s_x = s_y$, that is, with four independent parameters.

The potential location of the position parameters \mathbf{t} for the potential parameters \mathbf{A} of the linear transformation can be expressed as (Montiel et al., 2001)

$$\mathbf{t}(W, \mathbf{x}, \mathbf{A}) = \mathbf{x} - \mathbf{A}W. \quad (1)$$

This traces a curve in the parameter space, and after gathering all evidences for all edge pixels, the maxima of the accumulator array define the best values \mathbf{A}^* and \mathbf{t}^* which correspond to the transformations that map the model to the image. Note that from Eq. (1) it becomes clear that the parameters of the linear transformation can be seen as input, while the translation as output, of the detection.

The complexity of the extraction process is now completely reported on the transformation and is independent on the shape template. Indeed, until now, no assumptions have been made about W . Eq. (1) can be used for analytical as well as for non-analytical templates. In the remainder of the text, only non-analytical shapes will be dealt with. That is, the templates are given as a set of points as $W = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$.

When dealing with non-analytical objects, the shape information is encoded in a reference table (\mathbf{R} -table), in which the entry indices correspond to

the gradient orientation φ at the object boundary. For a given reference point $\bar{\mathbf{o}}$ (e.g. the gravity center or the center of the minimum bounding box to minimize the relative error), the vectors $\bar{\mathbf{r}}_i = \bar{\mathbf{o}} - \bar{\mathbf{x}}_i$ are stored as a function of φ (Ballard, 1981) for $i = 1, \dots, n$.

During the detection process, the elements of the accumulator array that are incremented are given by Aguado (2002)

$$\{(\mathbf{t}, \mathbf{A}) | \mathbf{t} = \mathbf{x} - \mathbf{A}\bar{\mathbf{x}}_i, |\varphi(\mathbf{x}) - \arg[\mathbf{A}(\cos(\bar{\mathbf{x}}_i), \sin(\bar{\mathbf{x}}_i))']| < \delta\varphi, \\ (i = 1, \dots, n)\} \quad \forall \mathbf{x} \in I,$$

where $\delta\varphi$ is an allowable angle error and I the input image. Under the assumption of weak affine transformations, this brute force voting GHT delivers a reliable estimation of the object position and a coarse approximation of its pose, which will be refined in a second step as described in the next subsection.

2.2. Adaptive generalized Hough transform for affine transformations

The parameters resulting from the previous step can be used as starting values of another HT for more complex transformations (e.g. affine). The parameter search can be therefore restricted to a small area around the previous solution, saving time and storage complexities. However, even in that case, if the transformation parameters are numerous and discretized with high accuracy, the array becomes intolerably large and the number of possible combinations too important for real applications.

Based on the idea of Illingworth and Kittler (1987), an adaptive coarse-to-fine approach is proposed. Unlike the HT which represents the whole parameter space by a quantized array, the adaptive strategy consists of accumulating the HT in a small size accumulator and using this information intelligently to redefine the parameter range so that only the interesting areas are investigated in greater detail. This procedure can be repeated until the parameters reach a pre-specified accuracy without increasing the storage space.

In order to keep interesting time performances, arrays of small size should be used. We chose a $3 \times 3 \times 3$ accumulator array for affine transformations. This is a reasonable compromise between speed and interpretation of the general space topology. Each cell of the array corresponds to a plane with image dimensions, which will be accumulated during the detection stage. After accumulation, the maximum is searched through all the planes. The cell corresponding to the plane having the highest peak returns the transformation parameters for a given accuracy, while the object position is given by the position of the maximum within this plane. The parameter range and accuracy for the next iteration can be re-estimated by analyzing the position of the maximum for each axis of the array according to the following rule:

- if the maximum was located on an exterior cell, then the accuracy is kept and the parameters are re-arranged such that this cell stands in the middle of the new array (\rightarrow location of the actual parameter, Fig. 1(a));
- else a new range is defined around the cell of maximum value a_{\max} as $[a_{\max} - \delta; a_{\max} + \delta]$ ($a = \{a_A, b_A, c_A, d_A\}$), which is then is quantized in three new bins of finer accuracy. The new accuracy is now given by $(2/3)\delta$, where δ is a small number decreasing over the time. For instance, δ can be updated as $\delta := 0.9 \cdot \delta$ every time that a new range is defined (\rightarrow accuracy improvement, Fig. 1(b)).




Fig. 1. Interpretation of the parameter space axes after accumulation. (a) The axis is centered around the cell of maximum accumulation. (b) If the maximum lies in the middle cell, the parameter accuracy is improved by reducing the axis range.

3. Accounting for local shape variability

The HT for the detection of shapes under weak affine transformations presented above does a good job even in presence of occluded shapes. However, the complete specification of the *exact* shape is required to achieve precise segmentation. For real life computer vision or medical applications, the target objects usually does not have exactly the same shape as the transformed template and local variations should be taken into account to improve the accuracy of the detection.

3.1. Local shape variability with active Hough transform

The HT is one of the most powerful methods for the detection of shapes under geometric transformations in images. However, until now, the HT has been assumed to be used for the recognition of a single shape only, for instance the middle shape of an object class. Natural objects are characterized by much more flexibility and a HT for natural shapes in supervised detection was introduced by Brejl and Sonka (2000), assuming that different prototypes of the object of interest are provided in a form of segmented example images. Let this training set be made out of N_s shapes W_1, \dots, W_{N_s} , ($W_j = (\bar{x}_{1j}, \dots, \bar{x}_{N_{ij}})$, $j = 1, \dots, N_s$), each of them being aligned to the mean shape and being composed of N_l corresponding samplings. The manually identification of these points may rapidly become very tedious and in real applications only few of them are identified by an expert user. The remaining of these points can be found by sampling the interval defined by two successive landmarks. Small sampling steps should be used to provide enough data for the building of the R-table.

Unlike the original GHT, all shapes of the training set will be encoded in the R-table, that is, all $N_s \cdot N_l$ points will now contribute to a new entry. This operation consists essentially of encoding the shape variability of an object class within the R-table. In opposition to (Brejl and Sonka, 2000), we argue that if an entry already exists, the new contribution should nevertheless be added. Indeed, the weight of a point, which appears several times

in the R-table, will be more important, since that contribution is more likely to happen.

3.2. Uncertainty regions for error compensation

Many sources of error influence the computation of the shape parameters \mathbf{A} and \mathbf{t} , so that not only one cell of the accumulator array is incremented, but rather a region around the expected point. This problem was already investigated by Ballard (1981) or Garrido et al. (1998). The latter proposed to associate to each R-table entry an uncertainty region, such that not just one single cell is incremented, but all the cells corresponding to that region. That is, an uncertainty region is given to each entry of the R-table whose size (width) depends, for instance, on an a priori deformation probability distribution. This solution is attractive since each uncertainty region can reflect the actual deformation at a given boundary point.

The uncertainty regions can be directly computed from the R-table of the active HT described above. Indeed, to each entry index of the R-table (i.e. discrete gradient orientation) corresponds a set position vectors relative to a reference point (Section 2.1). Considering these vectors as samples of a random distribution, it is now possible to define the uncertainty region as a density of probability which is directly estimated from these samples. The statistical modeling can be performed with the adaptive kernel method (Silverman, 1986), which automatically adapts the amount of smoothing according to the data, without user's knowledge. It becomes now clear, that during the building of the R-table, all $N_s \cdot N_l$ points of the training set should be considered, even if a similar entry has already been recorded.

3.3. Shape variance in the reconstruction step

The last issue to address is the reconstruction of the shape from the maximum of the accumulator array. The easiest solution would be to compute the back transformation of the mean shape, but it would not take into account neither the shape variability nor the contribution of the image.

To include shape variance in the reconstruction step, the corresponding landmarks of the training set can be used intelligently. Let $L_i = \{\bar{\mathbf{x}}_{i1}, \bar{\mathbf{x}}_{i2}, \dots, \bar{\mathbf{x}}_{iN_s}\}$ ($i = 1, \dots, N_l$) be the i th set of corresponding landmarks for the N_s shapes of the training set after back transformation according to the optimal parameters extracted from the maximum of the accumulator array. For $1 \leq i < N_l$, L_i can be seen as a set of N_s observations of a two-dimensional density of probability $p_{L_i}(x, y)$. Modeling the underlying probability density with the adaptive kernel method as above, we obtain $\hat{p}_{L_i}(x, y)$, an estimate of the distribution of the i th landmark points. To rely this information with the image content, the position of the boundary points can be computed by

$$\arg \max_{(x,y)} \{\hat{p}_{L_i}(x, y) \cdot p_B(x, y)\}, \quad (2)$$

where $p_B(x, y)$ is the boundary probability for a given image (Paragios, 2000). Maximizing Eq. (2) is equivalent to maximizing the joint probability for the position and boundary. It simultaneously guarantees likely shape variance along with good correspondence to image boundary. (For practical uses, $p_B(x, y)$ can also be seen as the gradient image rescaled between 0 and 1.)

4. Experimental results

In this section, two experiments are carried out on real medical images. The first experiment is concerned with the detection of the corpus callosum from 2D saggital magnetic resonance (MR) slices, while in the second experiment, the left ventricle is searched in 2D axial MR images. For an illustration of the searched anatomical structures, please refer to Fig. 2. Note that all images are of size 256×256 pixels and that for clarity, only zoomed parts will be shown in the next subsections.

4.1. Detection of corpus callosa

Fig. 3 illustrates the different steps of the HT for natural shapes under weak affine transformations and local geometric variations. Fig. 3(a) corre-




Fig. 2. Anatomical structures to extract in the experiments of Sections 4.1 and 4.2. (a) Corpus callosum, (b) left ventricle.

sponds to the mean shape of the object class, which was made of 12 instances of previously segmented corpus callosa. On each shape, 50 points were computed from 5 landmarks identified manually. Fig. 3(b)–(d) show the detection of different corpus callosa for the encoded shape class. The first and second rows correspond respectively to the detection results assuming a similitude, respectively an affine transformation, with rigid reconstruction from mean shape. The detection seems to be better for the second row, but it remains constrained to rigid transformations and cannot account for local deformations. Fig. 3(d) shows that in extreme cases, the generality of the affine transformation over the similitude can lead




Fig. 3. Detection of corpus callosa with the HT for shapes under weak affine transformations. (a) Mean shape of the object class. (b)–(d) Examples of detection for different patients. The first row shows the results for similarity transformations with rigid reconstruction. The second row shows the results for affine transformations with rigid reconstruction, whereas the third row corresponds to the results for affine transformations with reconstruction accounting for local shape variations.

to better results. The third row shows the results using shape variance in the reconstruction step.

Fig. 4 shows the translation parameter accumulator after the GHT for similitudes. From this figure, it can be seen that the peak corresponding to the object location is distinctly higher than the background noise, ensuring a good starting position for the adaptive HT for affine transformations.




Fig. 4. Translation parameter accumulator after the GHT for similarity transformations corresponding to column (a) of Fig. 3.




Fig. 5. Detection of left ventricles with the HT for shapes under weak affine transformations. (a) Mean shape of the object class. (b)–(d) Examples of detection for different patients. The first row shows the results for similarity transformations with rigid reconstruction. The second row shows the results for affine transformations with rigid reconstruction, whereas the third row corresponds to the results for affine transformations with reconstruction accounting for local shape variations.

The experiments were performed on a Linux station Pentium III, 1000 MHz with test images of size 256×256 pixels. The computation time required by the HT for the similarity transformation was about 3.5 s with a scale going from 0.8 to 1.5 quantized in 10 values and an orientation between -20.0° and 20.0° quantized in 20 values. These numerical values strongly depend on the application and were heuristically chosen, such that the range of possible transformations be covered. The computation time for the affine transformation depends on the number of iterations, which have to be performed until the pre-defined accuracy. In our examples, the initial δ was set to 0.15 and the program stops when the parameter precision reaches 0.01. A pass lasts 0.5 s and usually between 5 and 15 iterations are to be performed, leading to a maximum global time of $1.5 + 15 \cdot 0.5 = 9$ s.

This performance is usually acceptable for such applications.

4.2. Detection of left ventricles

Fig. 5 shows the results for the detection of the left ventricle. As previously, the rows of the figure compare the different steps of the proposed HT. The mean shape of the left ventricle class, composed of 12 previously segmented instances, is represented on Fig. 5(a) for 50 corresponding points computed from 5 manually identified landmarks per training image. The first row of Fig. 5 corresponds to the detection results assuming a similitude, whereas an affine transformation was assumed for the second and third rows. In the first two cases, the mean shape was used for the reconstruction, while the third row shows the results using shape variance in the reconstruction step (Section 3.3). Identical parameters were used as in the previous experiment, leading to similar time performances.

5. Concluding remarks

In this paper, an approach for the detection of non-analytic objects based on the Hough transform is presented. The usual formulation assumed a similitude mapping between the template and the shape to detect. Thanks to an adaptive implementation, the similitude can be extended to the more general affine transformation, while ensuring arbitrary accuracy and low storage requirement. The use of a brute force voting strategy makes it possible to detect the objects in complex environments.

Anatomical structures are soft and their deformations cannot only be characterized through geometric transformations. A significant improvement of the detection is achieved by incorporating local shape variability directly into the R-table. By this way, much more general shapes can be extracted. Moreover, unlike the original approach, we propose to reconstruct the shape from the peak of the accumulator array by combining both a priori shape variance and image fidelity.

Although robust detection can be achieved, large variations cannot be handled correctly.

Nevertheless, our algorithm can be seen to be the initialization of other segmentation algorithms like active contours. Combining both would be a step towards full automatic segmentation approaches.

References

- Aguado, A.S., Montiel, E., Nixon, M.S., 2000. Bias error analysis of the generalised Hough transform. *Journal of Mathematical Imaging and Vision* 12, 25–42.
- Aguado, A.S., Montiel, E., Nixon, M.S., 2002. Invariant characterisation of the Hough transform for pose estimation of arbitrary shapes. *Pattern Recognition* 35, 1083–1097.
- Ballard, D.H., 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13 (2), 111–122.
- Brejzl, M., Sonka, M., 2000. Object localization and border detection criteria design in edge-based image segmentation: Automated learning from examples. *IEEE Transactions on Medical Imaging* 19 (10), 973–985.
- Chakraborty, A., Staib, L.H., Duncan, J.S., 1996. Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Transactions on Medical Imaging* 15 (6), 859–870.
- Cootes, T., Hill, A., Taylor, C., Haslam, J., 1994. The use of active shape models for locating structures in medical images. *Image and Vision Computing* 12 (6), 355–366.
- Ecabert, O., Thiran, J.-P., 2002. Variational image segmentation by unifying region and boundary information. In: *ICPR*. pp. 885–888.
- Fung, P.-F., Lee, W.-S., King, I., 1996. Randomized generalized Hough transform for 2-D grayscale object detection. In: *ICPR*. Vol. 3. pp. 511–515.
- Garrido, A., Pérez De La Blanca, N., 1998. Physically-based active shape models: Initialization and optimization. *Pattern Recognition* 13 (8).
- Hough, P.V.C., 1962. Method and means for recognizing complex patterns. US Patent 3069654.
- Illingworth, J., Kittler, J., 1987. The adaptive Hough transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9 (5), 690–698.
- Kass, M., Witkin, A., Terzopoulos, D., 1998. Snakes: Active contours models. *International Journal of Computer Vision* 1, 321–331.
- Kimura, A., Watanabe, T., 2002. An extension of the generalized Hough transform to realize affine-invariant two-dimensional (2D) shape detection. In: *ICPR*. Quebec City, Canada.
- Leventon, M.E., Grimson, W.E., Faugeras, O., 2000. Statistical shape influence in geodesic active contours. In: *Proceedings of CVPR*. pp. 316–323.
- Montiel, E., Aguado, A.S., Nixon, M.S., 2001. Improving the Hough transform gathering process for affine transformations. *Pattern Recognition Letters* 22, 959–969.
- Paragios, N., 2000. Geodesic active regions and level set methods: Contributions and applications in artificial vision. Ph.D. thesis, University of Nice, Sophia Antipolis.
- Ser, P.-K., Siu, W.-C., 1995. A new generalized Hough transform for the detection of irregular objects. *Journal of Visual Communication and Image Representation* 6 (3), 256–264.
- Silverman, B.W., 1986. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall.
- Xu, L., Oja, E., 1993. Randomized Hough transform RHT: Basic mechanisms, algorithms, and computational complexities. *CVGIP: Image Understanding* 57 (2), 131–154.