

A Study on Anti-Geometric Diffusion for the Segmentation of Human Knee Cartilage

James Cheong and David Suter

Dept. of ECSE, PO Box 35, Monash University, Clayton 3800, Australia, 2004.

{james.cheong, [d.suter](mailto:d.suter@eng.monash.edu.au)}@eng.monash.edu.au

Abstract: Anisotropic diffusion was first introduced by Perona and Malik [1] for image smoothing and denoising. Since then, the field has matured and a better understanding of its properties and implementation has led to numerous applications for image processing. The objective of this study is to evaluate the effectiveness of anti-geometric diffusion as a method for segmenting knee cartilage from MRI scans. This report will give a detailed description of anti-geometric diffusion and investigate its use together with energy based region merging as a greyscale segmentation method proposed by Manay [2]. A description of the method as well as its implementation will be presented in this report. We will also display and discuss some of the results obtained from running Manay’s segmentation method on our library of knee MRI.

Key words: *Anti-Geometric Diffusion, Segmentation, Osteoarthritis, Cartilage,*

1 Introduction

According to the Arthritis Foundation of Australia, over 3.1 million people in Australia suffer from some form of arthritis. Arthritis means “joint inflammation” and is a serious chronic condition with no known cure. There is currently an urgent need to better understand the disease, as arthritis is the leading cause of disability among people over the age of 65. With an aging population and no known cure, the disease is expected to reach epidemic proportions in the near future.

Osteoarthritis is the most common form of arthritis and involves the gradual loss of articular (joint) cartilage^a. It affects about 14% of the adult population [3] and is most prevalent in the knee and hip joints. Recent studies have shown that the quantitative measurement of knee cartilage volume is an accurate and reproducible method for the measurement of osteoarthritis progression [4].

Current methods of cartilage volume measurement involve some form of manual segmentation carried out by a trained clinician. The key steps in the segmentation process involve delineating the cartilage and separating it from the surrounding tissues. Images of a patient's knee are obtained using magnetic imaging resonance (MRI) with fat-suppression to provide the best contrast between cartilage and bone. The scans obtained are usually greyscale images in the sagittal plane and consists typically of 60 images (slices) for each knee. Using some form of medical/imaging software, the clinician will visually inspect and identify the presence of cartilage on each image slice. If cartilage is present, the cartilage boundary is manually traced, see Figure 1. When all 60 slices have been processed, an interpolation function is applied, using the inter-slice distance (slice thickness) as a parameter to approximate the cartilage volume.

This manual process is laborious and can take up to 1 hour for each patient scan. It is also subject to the judgement of the clinician and requires significant experience and training to produce accurate and reproducible results. Thus, there is a strong demand for an improved automated method that will segment and measure the volume/surface area of articular cartilage in the human knee joint from MRI scans.

^a **Cartilage:** A tough, elastic, fibrous connective tissue found in various parts of the body, such as the joints, outer ear, and larynx. In the knee region, it is smooth and white, and lines the surface of the femur, tibia and patella within the joint.




Figure 1: MRI of the Knee in Sagittal Plane, Tibial Cartilage Manually Segmented by Clinician

Apart from the manual methods described, a number of semi-automated methods have been developed [5-9]. These usually involve a human operator initialising the procedure by setting a number of starting points, followed by an automated process that performs some image processing operations. Canny's edge detection filter, active contours (snakes), template matching, and cubic B-splines are some of the key components of the methods used.

The objective of this study is to explore the usefulness of anti-geometric diffusion as a segmentation method proposed by Manay [2]. Section 2 will introduce the concept of anti-geometric diffusion, followed by a description of the segmentation algorithm in Section 3. Section 4 will present the results of the segmentation algorithm on our database of MRI knee images. Finally Sections 5 and 6 will provide some discussion on the results and end with a conclusion.

2 Anti-Geometric Diffusion

Anisotropic diffusion was first introduced by Perona and Malik [1] for image smoothing and denoising. Their anisotropic diffusion method is formulated to encourage intraregion smoothing in preference to smoothing across region

boundaries. Most of the research on anisotropic diffusion over the years has been focused on the preservation of features in the image while denoising the image data. Anti-geometric diffusion is a form of anisotropic diffusion that goes against this trend by diffusing *across* image edges, in a direction orthogonal to the geometric heat flow. Geometric heat flow diffuses *along* image edges, thus preserving the edges while anti-geometric diffusion effectively spreads out the edge information. The method is thus termed anti-geometric because it is orthogonal to the geometric heat flow.

The advantage of smearing edge information is that it allows quick detection of features and their location within an image, thus enabling fast segmentation of the image. Image regions that lie nearby, but on opposite sides of a prominent edge are quickly distinguished.

Edge directions are usually related to the tangents of the isointensity contours (level curves), since the tangent direction of an isointensity contour is the direction perpendicular to the image gradient. Using the first derivatives of the image I_x and I_y , we can define η , the direction normal to a given point on a level curve and ξ , the tangential direction (see Figure 2), as:

$$\eta = \frac{(I_x, I_y)}{\sqrt{I_x^2 + I_y^2}} \quad \xi = \frac{(-I_y, I_x)}{\sqrt{I_x^2 + I_y^2}} \quad (1)$$




Figure 2: Normal and tangential directions of a point on a level curve

Image diffusion is based on partial differential equations. In the case of anti-geometric diffusion, the tangential diffusion is excluded and only the normal diffusion is applied. The result of diffusion in the normal direction is that the image edges are smeared. The anti-geometric diffusion equation is defined as:

$$\frac{\partial I}{\partial t} = \frac{I_x^2 I_{xx} + 2I_x I_y I_{xy} + I_y^2 I_{yy}}{I_x^2 + I_y^2} \quad (2)$$

3 Manay's Segmentation Method

Manay [2] introduces a segmentation method for greyscale images using anti-geometric diffusion together with energy based region merging. This method has not been tested on knee cartilage segmentation, but has proven to successfully segment cardiac MRI scans. The basic idea behind the segmentation method is to iterate between diffusion and region merging steps until a desired number of regions is reached, and the difference between the region merging steps is small. The desired number of regions is defined by the user. The diffusion step essentially breaks the image down into regions, or when applied to an undersegmented region, breaks it down into smaller regions to reveal small-scale features. The region merging step is carried out to compensate for the oversegmentation that occurs as a result of the diffusion process. When the segmentation process eventually stops, the resulting image should consist of the desired number of regions, and among these, one region will represent the femoral cartilage, and another region, the tibial cartilage.

This section will present the segmentation algorithm in detail, as well as some of the methods used for classification during the diffusion process.

3.1 Consistency Problem

One of the crucial steps involved while segmenting using diffusion is the classification of pixels. In the method proposed by Manay, classification is made during the diffusion process, instead of afterwards. A pixel is classified as soon as its classification criterion becomes unambiguous and this classification is maintained as the diffusion proceeds.

In the early stages of diffusion, only the intensity values of pixels near object boundaries change significantly. This is due to the nature of anti-geometric diffusion, which diffuses across image edges, and the fact that boundaries are defined by a sharp change in pixel intensity. As diffusion proceeds, the image edges are smeared even more and the global averaging effect enables classification of pixels further away from important image features like object boundaries. However, prolonged diffusion can result in diffused intensities near boundaries of smaller features changing from being brighter than the original intensities to darker than the original intensities (or vice-versa). By classifying a pixel when its classification criterion becomes unambiguous and maintaining this classification, consistency problems for pixels already classified are avoided. This also enables the diffusion process to be run for as long as is necessary to classify pixels far from region boundaries. The diffusion process is generally stopped when enough pixels have been classified.

3.2 Classification Criteria

A common classification criteria use during diffusion is to track a pixel's net intensity change. A single constant change in the pixel's diffusion intensity is used to decide when and how a pixel is classified. However, a more robust method would be to track whether a pixel's diffusion intensity is consistently increasing or decreasing (monotonic). The advantage of tracking monotonicity instead of absolute intensity changes is that the classification will not be sensitive to the magnitude of intensity changes, thus it will be more robust to noise.

For example, given an isolated bright pixel of noise on a dark background near the boundary of a large bright point, Figure 3; diffusion will initially result in the isolated pixel's diffusion intensity decreasing by a large amount as the noise is smoothed away. However, as diffusion continues, the pixel's intensity will increase due to the diffusion of the large bright region. Even though the latter increase in intensity may be smaller than the initial decrease in intensity, the initial decrease happens quickly while the latter increase may continue steadily for a long time. Therefore, classification by absolute intensity change will incorrectly classify the noise as part of

the bright region while classification by monotonicity will correctly classify the noise as part of the dark region.

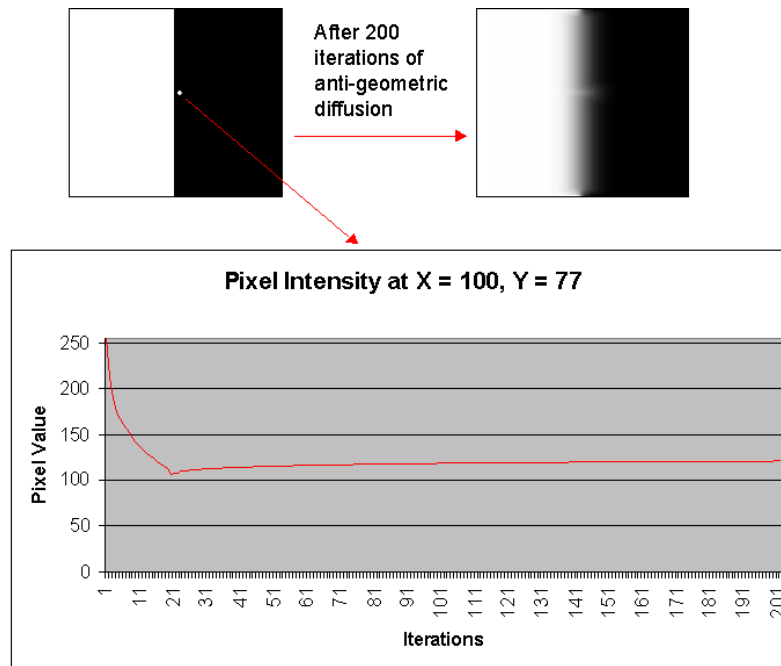


Figure 3: Plot of the intensity of the noise pixel as anti-geometric diffusion is applied

It is common to use a combination of monotonicity and absolute intensity change to increase the overall robustness of classification in the presence of noise.

3.3 The Algorithm

The segmentation method starts off by applying anti-geometric diffusion on the image for a short time, during which all pixels are classified into one of the three categories, 1) pixel intensities that diffused significantly and/or monotonically upward, 2) pixel intensities that diffused significantly and/or monotonically downward, and 3) pixels which did not diffuse significantly. Neighbouring pixels with the same classification are then grouped into connected regions, where each connected region R_i is assigned a unique label, i . Because anti-geometric diffusion acts as a region splitting operator, the original image is broken into smaller regions. Diffusion progresses until a preselected percentage of the pixels in the image are classified.

The next step involves energy based region merging suggested in [10, 11]. Statistics for each region are calculated and neighbouring pairs of regions R_i and R_j that yield the smallest increase in ΔE_{ij} are merged, where E_i is termed the squared error for R_i and ΔE_{ij} , the change in squared error when R_i and R_j are merged. E_i and ΔE_{ij} are defined as:

$$E_i = \sum_{R_i} (I - \mu_i)^2 \quad (3)$$

I = Image data, μ_i = Mean of region R_i .

$$\Delta E_{ij} = E_{ij} - E_i - E_j \quad \text{where} \quad E_{ij} = \sum_{R_{ij}} (I - \mu_{ij})^2 \quad (4)$$

$R_{ij} = R_i \cup R_j$, μ_{ij} = Mean of region R_{ij}

The number of regions is reduce by one each time two neighbouring pairs are successfully merged. This merging process is repeated until a desired number of regions are obtained or until the total squared error, E , reaches a predefined maximum limit. E is defined as the sum of the squared errors for all the regions in the image, i.e.:

$$E = \sum_i E_i \quad (5)$$

When the merging process stops, the region with the highest squared error is split using anti-geometric diffusion. This is because regions exhibiting large squared errors are likely to contain undersegmented regions. After this step of diffusion, the image is likely to be oversegmented and merging is again required to reduce the number of regions. This iterative process of splitting and merging is repeated until the total square error converges, i.e. $\frac{\Delta E}{E}$ becomes small.

3.4 Implementation

The implementation of anti-geometric diffusion suggested by Manay consists of 2 key steps per iteration of diffusion; the calculation of the new intensity value of each pixel

after diffusion for 1 timestep, and the classification of pixels according to their diffusion behaviour and the specified classification criterion.

The implementation of each anti-geometric diffusion step makes use of a standard forward Euler step:

$$I[i, j, t + \Delta t] = I[i, j, t] + \Delta t I_{\eta\eta}[i, j, t] \quad (6)$$

where

$$I_{\eta\eta} = \frac{A + 2B + C}{I_x^2[i, j, t] + I_y^2[i, j, t]}, \text{ where}$$

$$A = I_x^2[i, j, t] I_{xx}[i, j, t],$$

$$B = I_x[i, j, t] I_y[i, j, t] I_{xy}[i, j, t], \text{ and}$$

$$C = I_y^2[i, j, t] I_{yy}[i, j, t] \quad (7)$$

and

$$I_x[i, j, t] = \frac{I[i+1, j, t] - I[i-1, j, t]}{2\Delta x}$$

$$I_y[i, j, t] = \frac{I[i, j+1, t] - I[i, j-1, t]}{2\Delta y}$$

$$I_{xx}[i, j, t] = \frac{I[i+1, j, t] - 2I[i, j, t] + I[i-1, j, t]}{(\Delta x)^2}$$

$$I_{yy}[i, j, t] = \frac{I[i, j+1, t] - 2I[i, j, t] + I[i, j-1, t]}{(\Delta y)^2}$$

$$I_{xy}[i, j, t] = \frac{A - B}{4\Delta x \Delta y}, \text{ where}$$

$$A = I[i+1, j+1, t] + I[i-1, j-1, t], \text{ and}$$

$$B = I[i+1, j-1, t] + I[i-1, j+1, t] \quad (8)$$

The step size, Δt , should be chosen such that $\Delta t < 0.5(\Delta x)^2$. This is to maintain the stability condition of the forward Euler calculations. In cases where I_x and I_y both equal zero, $I_{\eta\eta}$ is reverted to the linear flow, i.e. $I_{\eta\eta} = I_{xx} + I_{yy}$. The diffusion is iterated until a predetermined percentage of the pixels in the image are classified. As mentioned before, after a pixel has been classified, its classification is maintained.

The classification criteria used by Manay involves a combination of time of monotonicity (Δt) and intensity. The classification criteria can be defined as:

$$\Delta Intensity + m\Delta t > Thresh \quad (9)$$

where the coefficient m and $Thresh$ are values determined by the user.

The pixels are then classified as either white or black depending on whether the pixel intensity diffused monotonically downward or upward respectively, or grey if the value calculated for the criteria is smaller than $Thresh$. It should be noted that the anti-geometric diffusion process has to be applied on the whole image initially, then on a specific region, the one with the largest squared error, after the region merging step. The approach Manay has taken is to have two anti-geometric diffusion functions, one that requires an image as an input parameter and applies diffusion on the whole image, and another function that takes a region (linked list of pixels) as an input parameter and applies anti-geometric diffusion only within the region.

The implementation of the region merging step involves the use of a graph data structure, where each node in the graph is use to represent a region in the image and a connection between two nodes indicate that the corresponding regions are adjacent. Each node will contain a linked list of all the pixels in the region as well as the following statistics about the region:

$$\begin{aligned} A_i & \rightarrow \text{area (number of pixels) within region } R_i \\ S_i & \rightarrow \text{sum of pixel intensities within region } R_i \\ Q_i & \rightarrow \text{sum of squared pixel intensities within region } R_i \\ \mu_i = S_i / A_i & \rightarrow \text{mean pixel intensity with region } R_i \\ E_i = Q_i - 2\mu_i S_i + \mu_i^2 A_i & \rightarrow \text{squared error for region } R_i \end{aligned}$$

The graph is constructed initially by performing a connected component analysis of the classified and unclassified pixels. After the graph has been constructed, we iterate through each node and calculate ΔE_{ij} for each of its adjacent regions in search of a

suitable merging pair. After all adjacent nodes have been analysed, the pair of regions with the smallest ΔE_{ij} are merged.

It should be noted that we do not need to scan through the image data repeatedly to compute ΔE_{ij} for each pair of adjacent regions R_i and R_j . Instead, after the initial graph has been constructed, the statistics of the proposed region $R_{ij} = R_i \cup R_j$ can be calculated using the statistics of R_i and R_j , and the following relationships:

$$\begin{aligned}
 A_{ij} &= A_i + A_j \\
 S_{ij} &= S_i + S_j \\
 Q_{ij} &= Q_i + Q_j \\
 \mu_{ij} &= S_{ij} / A_{ij} \\
 E_{ij} &= Q_{ij} - 2\mu_{ij}S_{ij} + \mu_{ij}^2 A_{ij} \\
 \Delta E_{ij} &= S_i\mu_i + S_j\mu_j - S_{ij}\mu_{ij}
 \end{aligned}$$

Region merging stops when the desired number of regions in the image is reached.

4 Results

This section will display some of the results of using Manay's segmentation method on our library of MRI knee images. As described in the previous section, there are a number of parameters that apply to the stop conditions for the diffusion and region merging process that the user has to determine beforehand. These are:

<i>Thresh</i>	– Threshold for classification criteria, see Equation 9.
<i>m</i>	– Coefficient of monotonicity for classification, see Equation 9.
<i>%Complete</i>	– Stop condition for diffusion, percentage of pixels classified.
<i>FNR</i>	– Stop condition for region merging, final number of regions.
<i>Tolerance</i>	– Stop condition for the segmentation algorithm, maximum tolerable change in total squared error of the image, E .

The following table shows the values of the five parameters set and the corresponding figures that were obtained. The parameters set for the first row are considered the “optimal” in a sense that these seem to produce the most consistent segmentation of

the cartilage across a range of input images, i.e. the tibial and femoral cartilage are clearly distinct regions with minimum noise fused to the regions. The other rows in the table have a single parameter varied from the “optimal” to highlight how each parameter affects the final output. It should be noted at this stage that the input images are cropped versions of the original 512x512 8 bit greyscale MRI images. This was done to save on processing time. These input images are 200x150 and contain the region of interest, namely the area surrounding the tibiofemoral joint.

Combination	Figures	Thresh	m	%Complete	FNR	Tolerance
1	Figure 4 to Figure 9	10	0.01	0.8	20	0.85
2	Figure 10 & Figure 11	10	0.01	0.8	40	0.85
3	Figure 12 & Figure 13	10	1	0.8	20	0.85
4	Figure 14 & Figure 15	5	0.01	0.8	20	0.85
5	Figure 16 & Figure 17	10	0.01	0.5	20	0.85

Table 1

Figure 4 and Figure 7 show the cropped versions of the MRI images MRIM1040 and MRIM1045 respectively. Figure 5 and Figure 8 show the result of running the first step of anti-geometric diffusion on the whole input image until 80% of the pixels have been classified. White indicates a pixel has diffused monotonically downward, black, monotonically upward, and grey that the pixel is still “unclassified”. The remaining images in this section show the final image that results from running Manay’s segmentation algorithm according to the different combination of parameters used.

Figure 6 and Figure 9 show good segmentation of the tibial cartilage using the “optimal” combination of parameters, but the femoral cartilage is fused with some background noise and surrounding tissue. In the case of Figure 9, because the tibial and femoral cartilage are in contact in some areas, they have been classified together as one region.




Figure 4: Cropped version of MRIM1040




Figure 5: MRIM1040 after 1st iteration of Anti-geometric Diffusion is stopped



Figure 6: Final 20 regions segmented for MRIM1040 using Combination 1




Figure 7: Cropped version of MRIM1045




Figure 8: MRIM1045 after 1st iteration of Anti-geometric Diffusion is stopped




Figure 9: Final 20 regions segmented for MRIM1045 using Combination 1



Figure 10: Final 40 regions segmented for MRIM1040 using Combination 2




Figure 11: Final 40 regions segmented for MRIM1045 using Combination 2

Figure 10 and Figure 11 show the result of choosing a high value for the final number of regions in the image. There is essentially not much difference in the shape of the cartilage detected compared with the “optimal” combination, however, some background noise has merged with the tibial cartilage in Figure 10. The additional regions detected don’t interfere with shape of the cartilage, but the resultant image is noisier because with a higher value for the final number of regions, the background noise is separated into different regions when they can actually be grouped together.




Figure 12: Final 20 regions segmented for MRIM1040 using Combination 3




Figure 13: Final 20 regions segmented for MRIM1045 using Combination 3

Figure 12 and Figure 13 show the result of choosing a large value for m , the coefficient for monotonicity. With a larger value for m , large chunks of background are merged with the femoral cartilage. The computation time is much faster than the “optimal” combination because the pixels require less diffusion timesteps before classification. The shapes of the tibial cartilage segmented are very similar to the “optimal” combination, but due to the large value for m , the effect of monotonicity on the classification procedure is weakened. This results in the surrounding tissues with relatively high pixel intensities merging with the cartilage region.




Figure 14: Final 20 regions segmented for MRIM1040 using Combination 4



Figure 15: Final 20 regions for segmented for MRIM1045 using Combination 4

Figure 14 and Figure 15 show the result of decreasing the value of *Thresh* to 5. In general, the shape of the cartilage regions are very similar to those obtained using the “optimal” combination. The computation time is also faster, as expected, because a pixel is classified more quickly due to the smaller threshold. However, some background regions have been separated into different regions when they should be merged together. This is the result of lowering the threshold and not allowing monotonicity to smooth out the noise.




Figure 16: Final 20 regions segmented for MRIM1040 using Combination 5



Figure 17: Final 20 regions segmented for MRIM1045 using Combination 5

Figure 16 and Figure 17 show the result of decreasing the *%Complete* variable to 50%. For image MRIM1040, a large portion of the background has fused with the femoral cartilage. With a lower value for *%Complete*, the first diffusion step quickly identifies the 2 cartilage regions. However, the region merging step incorrectly merges the femoral cartilage with a large region of unclassified background and this is maintained for the whole segmentation process. With image MRIM1045, the tibial cartilage has fuse with some noise.

There was no noticeable change in the output image if the tolerance was increased to 0.95, therefore no images were shown for this combination. It would not be sensible to have a low value for tolerance, because this would stop the algorithm prematurely.

5 Discussion

The results obtained using Manay's segmentation method were satisfactory. The algorithm performed well in images where the femoral and tibial cartilage are clearly distinguishable and not in contact. An initial problem encountered was deciding on the combination of input parameters to use. The "optimal" combination was obtained heuristically. This combination is optimal in the sense that we had the highest rate of successfully segmenting the cartilage into clearly distinct regions with minimal noise across a range of cropped knee images.

The main problem encountered with segmenting the cartilage involves the low contrast between cartilage and surrounding tissue in some areas of the joint surface, especially around joint contact areas, tendons and ligaments. The similar fat composition of these tissues results in similar greyscale values.

Another significant problem is that the range of pixel intensities that define a cartilage region can be quite large. In some slices, the range is 70 to 255 for an 8 bit greyscale image. Also, with images towards the middle of the 60 slices, the tibial cartilage thins out and this often results in there being more than one region of tibial cartilage. Figure 18 shows a good example of tibial cartilage with a wide range of pixel intensities and having more than one region. Figure 19 shows the ineffectiveness of Manay's method at successfully recognising the presence of cartilage when it has a large variation in intensity.




Figure 18: MRIM1058 with Tibial Cartilage Highlighted




Figure 19: Final 20 regions segmented for MRIM1058 using Combination 1

Some of the other problems encountered include cases where the tibial and femoral cartilages are touching, as in the case of MRIM1045. With such images, the two cartilages have been classified as a single region. The desired outcome of course is to have the two cartilages segmented into two separate regions, but because some parts of the cartilage are actually in contact, this is difficult to achieve using Manay's method, or any method based on pixel intensity values alone.

Another foreseeable problem is with severe cases of osteoarthritis, the femoral and tibial cartilage can be in contact and total cartilage volume is small, complicating detection and segmentation.

The implication of these problems is that Manay's method alone will not successfully segment cartilage for all the slices of a patient's MRI scan. It is difficult to rely simply on the pixel intensity to segment a cartilage region. Simple thresholding of the image will not separate cartilage from the surrounding tissue, and other methods like region growing from a seed pixel with fixed maximum and minimum threshold work only with limited success.

A human can easily detect the boundaries of the cartilage by taking into account the expected shape of the cartilage for the particular slice and the surrounding anatomy to infer the cartilage boundaries. With sufficient training and experience, a human can segment knee cartilage accurately and without too much difficulty, albeit a bit tedious.

Therefore, to successfully segment cartilage automatically, these factors of expected shape, surrounding anatomy and training have to be taken into account. For example, we could combine bone structure recognition with Manay's segmentation technique. When the tibia and femur is identified in an image, Manay's segmentation method is performed on the image. The classified regions resulting from the segmentation are then examined by the program to determine the likelihood that a region is the tibial or femoral cartilage based on parameters like proximity to the bone identified and shape of the region.

6 Conclusion

The use of anti-geometric diffusion as a means of segmenting human knee cartilage has been examined. The segmentation algorithm developed by Manay performs satisfactorily for most of the knee images in our library. It performs well when the femoral and tibial cartilages are clearly distinguishable and not in contact. The limitations of the method have been highlighted and discussed. Automatic segmentation of knee cartilage is a challenging problem and it seems unlikely that anti-geometric diffusion alone will solve the problem. It is suggested that Manay's segmentation method be used together with other techniques that make use of higher level knowledge such as anatomical structure to deal with the problem more effectively.

7 References

- [1] P. Perona, & Malik, J., "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 629-639, 1990.
- [2] S. Manay, & Yezzi, A., "Anti-Geometric Diffusion for Adaptive Thresholding and Fast Segmentation," *IEEE Transactions on Image Processing*, vol. 12, pp. 1310-1323, 2003.
- [3] M. D. Forman, Malamet, R., Kaplan, D., "A survey of osteoarthritis of the knee in the elderly," *The Journal of Rheumatology*, vol. 10, pp. 282-7, 1983.
- [4] F. Eckstein, Cicuttini, F., Raynauld, J. P., Waterton, J. C. & Peterfy, C., "Magnetic Resonance Imaging (MRI) of Articular Cartilage in Knee Osteoarthritis (OA): Morphological Assessment."
- [5] Z. A. Cohen, McCarthy, D. M., Kwak, S. D., Legrand, P., Fogarasi, F., Ciaccio, E. J., Ateshian, G. A., "Knee cartilage topography, thickness, and contact areas from MRI: in-vitro calibration and in-vivo measurements," *Osteoarthritis and Cartilage*, vol. 7, pp. 95-109, 1999.
- [6] A. A. Kshirsagar, Watson, P. J., Tyler, J. A., Hall, L. D., "Quantitation of articular cartilage dimensions by computer analysis of 3D MR images of human knee joints," pp. 753-756, 1997.
- [7] J. A. Lynch, Zaim, S., Zhao, J., Stork, A., Peterfy, C. G., & Genant, H. K., "Cartilage segmentation of 3D MRI scans of the osteoarthritic knee combining user knowledge and active contours," *Proceedings of SPIE*, vol. 3979, pp. 925-935, 2000.
- [8] T. Stammberger, Eckstein, F., Michaelis, M., Englmeier, K.-H., Reiser, M., "Interobserver reproducibility of quantitative cartilage measurements: comparison of B-spline snakes and manual segmentation.," *Magnetic Resonance Imaging*, vol. 17, pp. 1033-1042, 1999.
- [9] S. K. Warfield, Kaus, M., Jolesz, F. A. & Kikinis, R., "Adaptive, Template Moderated, Spatially Varying Statistical Classification.," *Medical Image Analysis*, vol. 4, pp. 43-55, 2000.
- [10] J. M. Morel, & Solimini, S., *Variational Methods in Image Segmentation*. Boston, MA: Birkhauser, 1995.
- [11] S. Zhu, & Yuille, A., "Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, pp. 884-900, 1996.