ANTI-GEOMETRIC DIFFUSION FOR ADAPTIVE THRESHOLDING AND SEGMENTATION

Siddharth Manay *Anthony Yezzi*

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332

ABSTRACT

In this paper, we present a novel adaptive thresholding technique based upon an anisotropic diffusion model, which may be referred to as the anti-geometric heat flow. In contrast to its more popular counterparts (such as the geometric heat flow) which diffuse parallel to image edges, this model diffuses perpendicular to image edges, yielding surfaces which are naturally suited for adaptive thresholding and segmentation. While it is possible to apply this diffusion for a fixed amount of time to detect features, we discuss how to detect features during the diffusion process, thus avoiding much of the arbitrariness associated with choosing a single scale (and makes the most notorious problem associated with anisotropic diffusion methods, namely "when do you stop?" a moot point). We will demonstrate the performance of this technique on both synthetic and real images, showing applications to thresholding written text and segmentation of medical images and scenes.

1. INTRODUCTION

Thresholding is one of the most widely used techniques in image processing and low-level vision, sometimes as an end goal and sometimes for preprocessing. The goal of thresholding is to create a binary image from a grayscale image, thus classifying the pixels into one of two categories (e.g. foreground and background). In this respect, thresholding is tantamount to segmentation when only two region classes are involved.

The most straight-forward approach to thresholding is to pick a fixed grayscale value (the threshold) and classify each image pixel by checking whether it lies above or below this value. Methods for judiciously choosing this value include: fit a pair of Gaussian curves to the histogram of a bimodal image, and choose the intensity value that minimizes the possibility of misclassification [3, 9]; construct discriminant functions, measures of class separability, and find the value that maximizes these functions [10]; or construct a criterion function related to the average pixel classification error rate, and find the minimum of this function [8].

Chow and Kaneko note in [3] that applying a single threshold to the entire image was not effective for images with spatially varying backgrounds. For such cases, they proposed dividing the image into subimages, finding thresholds for each subimage, and interpolating these local threshold values to construct a global thresholding surface. This idea can also be used in conjunction with the methods of Otsu or Kittler and Illingworth. The resulting *thresholding surface* or *adaptive threshold* gives rise to a spatially varying threshold which is used in the same manner as a single threshold to classify image pixels (by checking whether the value of a given pixel lies above or below the thresholding surface).

Yanowitz and Bruckstein [16] obtain an adaptive threshold by noting that pixel intensities near the transitions between foreground and background (edge pixels), in a smoothed image, serve as the best local thresholds. They locate such pixels by checking for large gradients and interpolate the grayscale values of these pixels to form the thresholding surface. Chan, Lam, and Zhu [2] outline a variational approach for obtaining an adaptive thresholding surface. Their method is accurate near edges, but the surface is very sensitive to local grayscale variations in regions far away from edges. This can produce false classifications in these regions.

A number of adaptive thresholding methods have been proposed which do *not* directly involve a thresholding surface. Many of these techniques approach the thresholding problem as a special case of segmentation (with two region classes). Intensity Gradient Based Thresholding [12] locates edge pixels in the image based upon the intensity gradient, classifies those pixels as foreground, and uses region growing techniques to classify the remaining pixels. For a fuller discussion of segmentation-based thresholding methods (and a more extensive list of references) see [4, 11]. Finally, we note that much work in thresholding has been done for the specific application of document binarization. We refer the interested reader to [13, 14] for a complete discussion of this application.

The intuition behind our model begins with the observation of Yanowitz and Bruckstein [16] that information about the best local threshold is to be found near image edges (transitions between foreground and background) *on a smoothed version of the image.* The smoothing step is crucial since the extremal greyscale values one encounters near a perfect edge will form very poor thresholds when compared to the "average" greyscale value encountered in the middle of a smoothed edge. The question is: "What is the best way to smooth?"

It is well understood that there is a close relationship between diffusion and other image smoothing techniques such as convolution. Accordingly, we will exploit a diffusion process which we believe is best suited to the goals of adaptive thresholding in order to capture the discriminative information embedded within the image edges. In particular, we will utilize an anisotropic diffusion model which, in contrast to most popular anisotropic diffusion models, diffuses specifically *across* image edges as opposed to *along* image edges. The latter behavior is desired for image denoising but the former behavior is better suited to adaptive thresholding since it has the effect of "spreading out" the edge information as far as possible. This will allow us to classify pixels near

0-7803-6725-1/01/$10.00 ©2001 IEEE

Authorized licensed use limited to: Georgia Institute of Technology. Downloaded on October 6, 2008 at 21:19 from IEEE Xplore. Restrictions apply.

image edges quickly and accurately. Then, as in Intensity Gradient Based Thresholding [12], we may extend these classifications to the remaining pixels (via interpolation or region growing, for example) or we may diffuse even further to classify additional pixels.

2. ANTI-GEOMETRIC DIFFUSION

In this section we discuss how diffusion may be used for adaptive thresholding and propose a diffusion model which seems to be ideally suited for these techniques.

A standard adaptive thresholding technique is to form a *thresholding surface* over the domain of an image and then classify image pixels based upon whether their values lie above or below this surface. A straight-forward method for constructing the thresholding surface is simply to blur the image with a Gaussian low pass filter. This is equivalent to diffusing the image via the linear heat equation, giving rise to a family of thresholding surfaces which comprise a well known scale space [15]. Consequently, choosing any particular threshold surface from this continuum imposes a certain scale on the features that are captured in the resulting binarized image. Near an image edge, a local average (fine scale) will yield an effective threshold, whereas away from an edge, a more global average (coarse scale) is necessary. It is not always clear which scale to choose.

Before addressing the ambiguity of scale (see Section 3), we note that it is natural to generalize this basic thresholding algorithm by using anisotropic diffusion. Anisotropic diffusion possesses the advantage of allowing local directional control of the diffusion process. This is particularly important where salient image features are concerned. When the preservation of sharp edges is important (as in image denoising), it is natural to consider models which diffuse along, but not across, the edge directions.

Typically, edge directions are related to the tangents of the isointensity contours (level curves or level sets) of an image I. Let η denote the direction normal to the level curve through a given point (the gradient direction), and let ξ denote the tangent direction.

Since η and ξ constitute orthogonal directions, we may express the rotationally invariant Laplacian operator as the sum of the second order spatial derivatives $I_{\eta\eta}$ and $I_{\xi\xi}$ in these directions and write the linear heat equation as

$$\frac{\partial I}{\partial t} = \nabla \cdot (\nabla I) = I_{\xi\xi} + I_{\eta\eta}.$$

Omitting the normal diffusion while keeping the tangential diffusion yields the well known *geometric heat flow*, which diffuses along the boundaries of image features, but not across them. It derives its name from the fact that, under this flow, the level curves of the image evolve in the normal direction in proportion to their curvature. This model is well known for its ability to denoise images while maintaining sharp edges and is therefore widely used for image enhancement and smoothing. For a more extensive discussion of the many properties of this flow see [1, 5, 6, 7].

The very property which makes the *geometric heat flow* powerful for image denoising (i.e. its ability to preserve edges in the image) makes it a poor flow for constructing adaptive thresholding surfaces, in which case we actually want to *smear* the image edges. If, instead, we omit the tangential diffusion and keep the normal diffusion, we obtain the complementary diffusion model, which we will refer to as the *anti-geometric heat flow*

Fig. 1. *Anti-Geometric, Geometric, and Linear Diffusion* TOP ROW: Ellipse blurred by (left-to-right) anti-geometric, geometric, and linear heat flows with equal diffusion times. BOTTOM ROW: Regions clearly identified as interior (black) or exterior (white) to the ellipse by comparing the original and diffused images. Grey regions indicate pixels whose original and diffused intensities differ by less than a small, uniformly fixed value.

$$\frac{\partial I}{\partial t} = \frac{I_x^2 I_{xx} + 2I_x I_y I_{xy} + I_y^2 I_{yy}}{I_x^2 + I_y^2}, \qquad (1)$$

in which diffusion occurs deliberately across the boundaries of image features. This is precisely what we want to occur when constructing an adaptive thresholding surface. Furthermore, by omitting the tangential *geometric* component of the diffusion, we avoid the shrinkage of the isointensity contours that occurs in both the geometric and the linear heat flow. Intuitively, the family of isointensity contours which run through a given edge are spread apart, while the shapes of those that remain near the original edge location are less distorted than they would be under the curvature-based shrinkage that would be induced by the discarded tangential diffusion.

In Figure 1 we illustrate the effects of the anti-geometric, the geometric, and the linear heat flows for generating adaptive thresholding surfaces for a synthetic image of an ellipse. The smearing of the ellipse edges induced by the anti-geometric flow is uniform in all directions. The resulting thresholding surface allows clear classification of all the pixels in a neighborhood of the boundary of the curve, regardless of the curvature of the ellipse. In the case of the geometric heat flow, there is very little blurring of edges. However, the ellipse shrinks, especially at the high-curvatures "tips." The thresholding surface formed with this flow can only classify the regions just inside these "tips," and nothing else. The linear heat flow represents a compromise, since it contains both geometric and anti-geometric components. As in the anti-geometric flow, information flows away from the boundary, and the corresponding thresholding surface provides clear classifications both inside and outside the boundary. However, due to the mitigating effect of the geometric component, the information flow near the "tips" of the ellipse is not as strong as in the anti-geometric case, and pixels outside these "tips" are not classified unless they are very close to the boundary.

Authorized licensed use limited to: Georgia Institute of Technology. Downloaded on October 6, 2008 at 21:19 from IEEE Xplore. Restrictions apply.

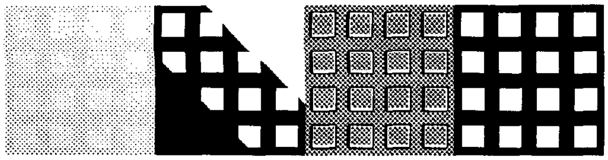

Fig. 2. Synthetic shaded image [left] thresholded using a fixed threshold [left center] and via anti-geometric diffusion [right center and right].

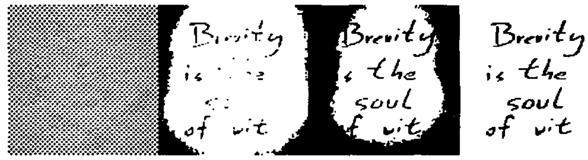

Fig. 3. Text image [left] thresholded using two different fixed thresholds [middle] and via anti-geometric diffusion.

3. THRESHOLDING VIA DIFFUSION

In this section we discuss how to use the anti-geometric diffusion process described in Section 2 for adaptive thresholding. One simple idea is simply to use an anti-geometrically diffused version of the original image as a thresholding surface. This approach, however, suffers from an arbitrary choice of diffusion time (which relates to a particular scale).

Rather than using a traditional thresholding surface (obtained using anti-geometric diffusion, however), we instead seek to classify pixels *during the diffusion process*. The intuition behind this is clarified by Figure 1 which shows that the "clearly classifiable" pixels identified by comparing an original and diffused image are located near region boundaries. As the diffusion proceeds, pixels further away from the boundaries become classifiable as well. Unfortunately, if we wait long enough for diffused intensities of pixels far away from region boundaries to differ from their original intensities enough to yield an unambiguous classification, diffused intensities near boundaries of smaller features may switch from being brighter than their original intensities to darker than the original intensities (or vice-versa) due to the more global averaging effect of prolonged diffusion. In other words, the global effect of prolonged diffusion helps us in classify distant pixels but could hurt us in maintaining a consistent classification of nearby pixels. If, however, we classify a given pixel *as soon as* its classification becomes unambiguous (i.e. the diffused and original intensities differ significantly) and *maintain* this classification as the diffusion proceeds, then we may run the diffusion as long as necessary to classify pixels far away from region boundaries without worrying about consistency problems for pixels that have already been classified. In this manner, we are no longer utilizing a single thresholding surface, but an *entire family* of thresholding surfaces generating by our anisotropic diffusion model.

This method is effective because pixels in regions with high detail (i.e. high spatial variance) change intensity relatively quickly during diffusion; their intensities can be unambiguously classified in short periods of time, and are therefore thresholded at a fine scale. Pixels in low-detail regions change intensity slowly, are thresholded at a much later time, and therefore are thresholded with a surface corresponding to a coarse scale.

One advantage of this method is that pixels near edges are quickly classified during the anisotropic diffusion process without the need for explicit edge detection prior to thresholding. A slight drawback is that pixels far away from region boundaries might only be classified after an extensive amount of diffusion. On the other hand, long runs of diffusion are not necessary to classify such pixels. If an image is truly bimodal, the pixels near the boundary of each region will be quickly and uniformly classified in the beginning of the diffusion process. The remaining pixels in this region will then be classified correctly by merely extending the classification from its boundary to fill in the remaining unclassified pixels.

This method is illustrated on a synthetic image in Figure 2. The first two images show a shaded image of 16 squares followed by a thresholded version using a fixed threshold. The third image demonstrates the partial classification via anti-geometric diffusion of the image on the left, followed by an extension of these classifications to the remaining pixels in the last image.

A pixel's net intensity change is not the only criterion that may be used to decide when to classify it during the diffusion process. In cases where a single value will not suffice near all image edges (the danger of merely choosing a tiny jump that is small enough to work for even the faintest image edges is sensitivity to noise) a better criterion is to check whether a pixel's diffusing intensity value is consistently increasing or consistently decreasing. Another alternative is to use a combination of intensity change, monotonicity requirements, and other criteria.

4. APPLICATIONS AND SIMULATIONS

In this section, we demonstrate the use of anti-geometric diffusion thresholding on several classes of images.

The image of handwritten text shown on the left in Figure 3 is extremely low-contrast with non-uniform illumination. Fixed thresholding is unable to capture all of the text, as shown in the two middle images. The far right image demonstrates the use of anti-geometric diffusion.

A nice feature of this diffusion approach is that we are not restricted to thresholding with just two regions. This thresholding model, when run on a multi-modal image, will differentiate pixels near edges quickly and accurately in relation to the edge, but might oversegment regions that are bordered by regions with brighter pixels and also by regions with darker pixels. However, an oversegmented image can often be a useful preprocessed input into a higher level segmentation algorithm. This is demonstrated in Figures 4 and 5 in which a cardiac MR image and a scene image (top left) are thresholded by running anti-geometric diffusion to classify between 40% and 50% of the pixels in the image. Then, in the final image (top right), the resulting classified and unclassified regions are "colored" with the mean intensity value of each corresponding region in the original image. We are investigating post-processing procedures to merge these segmented regions according to various criteria; these methods will combat the "flaws" in Figure 5.

Authorized licensed use limited to: Georgia Institute of Technology. Downloaded on October 6, 2008 at 21:19 from IEEE Xplore. Restrictions apply.

Fig. 4. Cardiac MRI [top left] segmented with anti-geometric diffusion and mean-coloring [top right]. For comparison, the image is segmented with two flat thresholds [middle and lower left] and mean-colored [middle and lower right]. (Images may be severely corrupted by dithering; printing this page on a laser printer is recommended.)

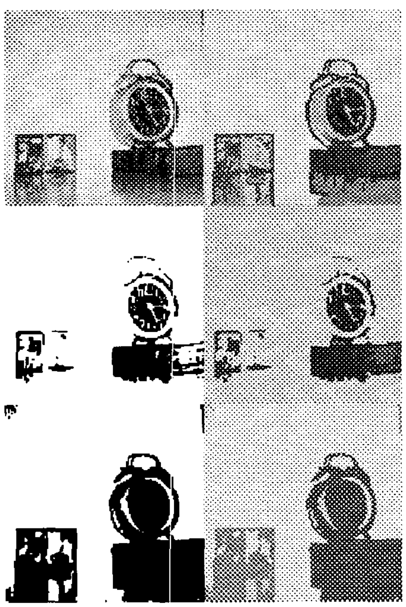

Fig. 5. Scene image [top left] segmented with anti-geometric diffusion and mean-coloring [top right]. For comparison, the image is segmented with two flat thresholds [middle and lower left] and mean-colored [middle and lower right]. (Images may be severely corrupted by dithering; printing this page on a laser printer is recommended.)

5. CONCLUSION

In this paper we have outlined a novel method for adaptive thresholding using an anti-geometric diffusion model to classify pixels in a greyscale image. In contrast to traditional approaches to adaptive thresholding, our formulation does not depend upon a single thresholding surface but an entire family of thresholding surfaces generated by a diffusing image. Pixels are classified *during* the diffusion process in this technique. We have shown how this technique applies not only to problems in which thresholding is the end-goal but also to preprocessing for segmentation.

6. REFERENCES

[1] L. Alvarez, P. L. Lions, and J. M. Morel, "Image Selective Smoothing and Edge Detection by Nonlinear Diffusion. II" *SIAM J. Numer. Anal.*, Vol. 29, pp. 845-866, 1992.

[2] F. H. Y. Chan, F. K. Lam, and H. Zhu, "Adaptive Thresholding by Variational Method," *Image Processing, IEEE Trans. on*, Vol. 7, no. 3, pp. 468-473, March 1998.

[3] C. K. Chow and T. Kaneko, "Automatic Boundary Detection of Left Ventricle from Cineangiograms," *Computer Biomed. Res.*, Vol. 5, pp. 338-410, 1972.

[4] R. M. Haralick and L. B. Shaprio, "Image Segmentation Techniques," *Computer Vision, Graphics, and Image Processing*, Vol. 29. pp. 100-132, 1983.

[5] B. B. Kimia and K. Siddiqi, "Geometric Heat Equation and Nonlinear Diffusion of Shapes and Images," *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1994.

[6] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "Shapes, Shocks, and Deformations I," *Int. Journal of Computer Vision*, Vol. 15, pp. 189-224, 1995.

[7] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "On the Evolution of Curves via a Function of Curvature, I the Classical Case," *Journal of Mathematics, Analysis, and Applications*, Vol. 163, pp. 438-458, 1992.

[8] J. Kittler and J. Illingworth, "Minimum Error Thresholding," *Pattern Recognition*, Vol. 19, no. 1, pp. 41-47, 1986.

[9] Y. Nakagawa and A. Rosenfeld, "Some Experiments on Variable Thresholding," *Pat. Rec.*, Vol. 11, no. 3, pp 191-204, 1979.

[10] N. Otsu, "A Threshold Selection Method from Gray Level Histogram," *IEEE Trans. Systems, Man, and Cybernetics*, Vol. SMC-8, pp. 62-66, 1979.

[11] N. R. Pal and S. K. Pal, "A Review on Image Segmentation Techniques," *Pattern Recognition*, Vol. 26 pp. 1227-1249, 1993.

[12] J. R. Parker, "Gray Level Thresholding in Badly Illuminated Images," *Pattern Analysis and Machine Intellegence*, Vol. 13, no. 8, pp. 813-819, Aug 1991.

[13] J. Sauvola, T. Seppänen, S. Haapakoski, and M. Pietikäinen, "Adaptive Document Binarization," *Document Analysis and Recognition, 1997, Proceedings of the Fourth International Conference on*, Vol. 1, pp 147-152, 1997.

[14] O. D. Trier and T. Taxt, "Evaluation of Binarization Methods for Document Images," *Pattern Analysis and Machine Intellegence, IEEE Transactions on*, Vol. 17, no. 3, pp. 312-315, March 1995

[15] A. P. Witkin, "Scale space filtering," *Int. Conf. on Artificial Intel.*, pp. 1019–1023, Karlsruhe, Germany, 1983.

[16] S. D. Yanowitz and A. M. Bruckstein, "A New Method for Image Segmentation," *Comp. Vision, Graphics, and Image Proc.*, Vol 46, no. 1, pp. 82-95, 1989.

Authorized licensed use limited to: Georgia Institute of Technology. Downloaded on October 6, 2008 at 21:19 from IEEE Xplore. Restrictions apply.