Grey Scale Skeletonisation with Curvature Sensitive Noise Damping

Huaijun Qiu and Edwin R. Hancock

Department of Computer Science, University of York, York, YO10 5DD, UK
{jun,erh}@cs.york.ac.uk

Abstract. This paper describes a method for curvature dependent skeletonisation in grey-scale images. We commence from a magnetostatic analogy, where the tangential edge flow is intepretted as a current. A vector potential is constructed by integrating the current weighted by inverse distance over the image plane. The skeleton corresponds to the location of valley lines in the vector potential. To damp noise effects we damp the current with an exponential function of the local curvature. In addition, we describe a number of postprocessing steps that can be used to improve the quality of the detected skeletons. In the end, we compare the effects of two alternative ways for noise damping.

1 Introduction

Skeletal abstractions have been used to great effect in the representation and recognition of both 2D and 3D shapes. Some of the earliest work was performed by Blum [2], who showed how the skeleton could be used for the morphological analysis of biological forms. Most of the existing work of skeletonisation has focused on binary valued objects [1, 7–9, 12]. Here a number of methods have been investigated including the medial axis transform [1], the chordal axis transform [7], and the grassfire transform [2], More recent work has focused on the analysis of the skeleton as the singularities in the eikonal equation for inward boundary motion [8]. An analysis of this system using the Hamilton-Jacobi equations of classical mechanics has shown how the skeleton can be detected using the divergence of the distance map for the object boundary [9]. Recently, Torsello and Hancock [12] have shown how the Hamilton-Jacobi skeleton can be improved by modifying the eikonal equation to take into account curvature effects.

In addition there has been a limited effort directed at the analysis of grey-scale objects. For instance, Tari, Shah and Pien [10] have proposed a linear diffusion equation to smooth out noise and extract skeletons directly from the grey scale images. Tek et al [11] have shown that an orientation sensitive distance propagation function can be used to extract symmetries from fragmented contours by labelling skeletal points according to whether or not they represent the collision of consistently oriented boundary fronts. Cross and Hancock [3] have appealed to a magnetostatic analogy in which the edge tangent flow is regarded as a current density on the image plane. The differential structure of the resulting vector potential can be used to characterise symmetry lines, boundary-edges

A. Fred et al. (Eds.): SSPR&SPR 2004, LNCS 3138, pp. 461–469, 2004.

and corners [5]. By sampling the vector potential at various heights above the image plane, a scale-space representation is induced. One of the shortcomings with this method is that like the binary skeleton, there is no way of moderating the effects of high curvature effects, other than by smoothing. This can prove time consuming and has the effect of removing genuine boundary structure.

The aim in this paper is to return to the magnetostatic analogy, and to incorporate curvature dependent damping of the current, i.e. the edge tangent flow. This has the effect of controlling boundary noise and improving the shapeliness of the skeleton. To improve the postprocessing of the vector potential, we apply hysteresis linking to the candidate skeleton points. In the end, we demonstrate how can we control the degree of smoothing of high curvature boundary features.

2 Image Representation Using Vector Potential

We commence by convolving the raw image I with a Gaussian kernel of width σ. The kernel takes the following form

$$G_\sigma(x,y) = \frac{1}{2\pi\sigma^2}\exp\left[-\frac{x^2+y^2}{2\sigma^2}\right] \tag{1}$$

With the filtered image to hand, the Canny edge map is recovered by computing the gradient

$$\underline{E} = \nabla G_\sigma * I \tag{2}$$

In order to compute a vector field representation of the edge-map, we will need to introduce an auxiliary z dimension to the original $x-y$ co-ordinate system of the plane image. In this augmented co-ordinate system, the components of the edge-map are confined to the image plane. In other words, the edge-vector at the point $(x,y,0)$ on the input image plane is given by

$$\underline{E}(x,y,0) = \begin{pmatrix} \frac{\partial G_\sigma * I(x,y)}{\partial x} \\ \frac{\partial G_\sigma * I(x,y)}{\partial y} \\ 0 \end{pmatrix} \tag{3}$$

For an ideal step-edge, the resulting image gradient will be directed along the boundary normal. In order to pursue our magneto-static analogy we would like to interpret the raw edge responses as elementary currents which flow around the boundaries and give rise to a vector potential. In other words, we would like to organise the elementary currents so that they are tangential to the boundaries of physical objects. Accordingly, we re-direct the edge-vectors to that they are tangential to the original planar shape by computing the cross-product with the normal to the image plane $\hat{z} = (0,0,1)^T$. The elementary current-vector at the point $(x,y,0)$ on the input image plane is defined to be

$$\underline{j}(x,y,0) = \hat{z} \wedge \nabla G_\sigma * I(x,y) \tag{4}$$

The key idea underlying the image representation is to characterise edges and symmetry lines using a vector potential. Edges corresponded to locations where

the elementary current re-enforce one-another. In other words, the boundaries are identified as local maxima of the vector potential. Symmetry points are those at which there is cancellation between diametrically opposed elementary currents. Axes of symmetry are lines of local minimum in the vector potential. At the level of fine detail, intensity ridges or ravines (lines) give rise to local symmetry axes.

According to magneto-statics, the vector-potential associated with a field of elementary currents is found by integrating over volume and weighting the contributing currents according to inverse distance. In other words, the vector potential at the point $\underline{r} = (x, y, z)^T$ in the augmented space in which the original image plane is embedded is

$$\underline{A}(x, y, z) = \mu \int_{V'} \frac{\underline{j}(x', y', z')}{|\underline{r} - \underline{r}'|} dV' \tag{5}$$

where $\underline{r}' = (x', y', z')^T$ and μ is the permeability constant which we set equal to unity. Since the contributing currents are distributed only on the image plane, the volume integral reduces to an area integral over the image plane. As a result, the components of the vector potential are as follows

$$\underline{A}(x, y, z) = \begin{pmatrix} -\int\int \frac{\partial G_\sigma * I(x', y')}{\partial y'} \frac{1}{\sqrt{(x-x')^2+(y-y')^2+z^2}} dx' dy' \\ \int\int \frac{\partial G_\sigma * I(x', y')}{\partial x'} \frac{1}{\sqrt{(x-x')^2+(y-y')^2+z^2}} dx' dy' \\ 0 \end{pmatrix} \tag{6}$$

The structure of the vector-potential deserves further comment. In the first instance, the components are confined to the $x - y$ plane for all values of the auxiliary co-ordinate z. However, as we move away from the image plane the role of this auxiliary dimension is to average the generating currents over an increasingly large area of the original image plane. In other words, if we sample the vector-potential for various $x - y$ planes at increasing height above the image plane, we induce a scale-space representation. We exploit this property to produce a fine-to-coarse image representation as we sample the vector potential at increasing heights above the physical image plane.

3 Curvature Estimation

In this section we consider how to incorporate curvature dependent smoothing into the vector potential to control the effects of high curvature boundary noise. To measure the curvature we use the method developed by Harris [4] which is itself an extension of Moravec's [6] corner detector. We commence by approximating the Hessian matrix $H = \nabla\nabla^T I$ by the matrix $E = (\nabla I)(\nabla I)^T$, i.e.

$$E = \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \tag{7}$$

where I_x and I_y are the first-derivatives of the image I in the x and y directions. Suppose that α and β are the eigenvalues of the symmetric matrix E. The

eigenvalues are proportional to the principal curvatures of the image intensity function, and the product of eigenvalues $\alpha\beta$ is hence proportional to the Gaussian curvature, while the sum of eigenvalues $\alpha + \beta$ is proportional to the mean curvature. According to Harris, the curvature response is the weighted sum of mean and Gaussian curvatures $R = \alpha\beta - k(\alpha + \beta)^2$. The quantity R can be used to characterise image features using the following tests

$$\begin{array}{l} R > 0 \; detected \; corners \\ R < 0 \; detected \quad edges \\ R \approx 0 \quad flat \qquad area \end{array}$$

We use the curvature measure to damp the current density. We adopt a model in which the effect of the current decays exponentially with the curvature of the image. The modified current density is

$$\underline{j}(x, y, 0) = \left[\hat{z} \wedge \nabla G_\sigma * I(x, y)\right] e^{-K \cdot R_{x,y}}$$

where K is a constant. Hence, the contribution from straight boundary segments is enhanced and the contributions from corners or places where the edge direction changes rapidly are suppressed.

By varying the constant K, we can control the degree of smoothing of high curvature boundary features. When K equals zero , then the original current density is recovered. When K is increased, the amount of curvature suppression is increased. As we will demonstrate later, the method of smoothing effect of boundary noise does not blur away genuine skeleton structures.

With the current density to hand, the vector potential is computed by performing the volume integration in Equation 6.

4 Skeletonisation and Noise Elimination

According to Cross and Hancock's [3] representation of image structure, symmetry lines follow the local minima of the vector potential and edge contours follow the local maxima. When viewed from the perspective of the differential structure of the vector potential, symmetry lines are locations where the component of the curl in the image plane vanish, i.e. $\hat{z} \wedge \nabla \wedge A(x, y, z) = 0$; edges are locations where the transverse component of the divergence vanishes, i.e. $\nabla \cdot (\hat{z} \wedge A(x, y, z)) = 0$. The main problem with this method for feature localisation is that it is subject to noise. In their work, Cross and Hancock fitted a prism surface to localise symmetry lines. Here we adopt a more sophisticated approach.

We commence by applying hysteresis linking to the edge (ridge) and symmetry (valley) lines in the vector potential to improve connectivity. For the edge (ridge) lines we perform connected components analysis. We dismember the web of ridge and valley lines at the locations of T-junctions. We then note the whereabouts of closed edge (ridge) lines. We remove those symmetry (valley) lines

that fall outside the closed edge contours, and retain only those that remain in the interior.

We illustrate the steps of our algorithm. In Figure 1(a) we show the raw image used in our study. Figures 1(b) and 1(c) show the principal curvatures. The maximum curvature is large along the edge contours of the image while the minimum curvature is large only at the locations of corners and noise. Figure 1(d) shows the magnitude $|\underline{j}|$ of the damped current density. White points correspond to locations where there is strongest damping, while the darker points correspond to strong current density. In Figure 1(e) we show the magnitude of the vector potential. Turning our attention to the postprocessing of the vector potential, the detected ridge and symmetry lines are shown in Figure 1(f). Figure 1(g) shows the result of applying hysteresis linking to the symmetry (valley) lines. After connected components analysis is performed, and the external symmetry lines have been removed then the resulting skeleton is shown in Figure 1(h). If the skeleton is dismembered at T-junctions, and branches with weal response are removed the result shown in Figure 1(i) is obtained. Finally, we fit straight lines to the detected symmetry lines and merge lines that are nearly parallel and close to each other, to obtain the result shown in Figure 1(j).

5 Experiment

In this section, we provide some experimental evaluation of our noise-damped vector potential representation. The experimental work is divided into two parts. We commence with some examples on synthetic images to illustrate the effect of K in damping the noises compared with the similar effect from increasing height. Next we furnish some real-world examples.

On the first row of Figure 2, we show three synthetic images with small spikes on the boundary and with a grey scale gradient on the interior. From the second row, we show in turn the magnitude of the vector potential displayed as a surface plot D, and, the correspondent detected ridges and ravines in black and white. For the first three rows of plots of vector potential, we increase the parameter K from 0 to 1, and then to 1.5. As we increase K, the effect is to damp-out noise while retaining the detail of the skeleton. As explained earlier, we can endow our image representation with a scale-space dimension by sampling the vector potential at increasing heights above the image plane. This is shown on the last two rows of vector potential plot by increasing the sampling height as we descend the columns. The effect of increasing the sampling height is also to smooth away noise, but this is as the expense of detail in the detected skeleton.

In Figure 3, we show some real-world images. There are four sets of images, grouped vertically for the first three with another one left at the bottom. For the vertical groups, the top panel shows the original image, the second panel shows the initial output, the third panel and the bottom panel show the output smoothed by increasing K and height respectively. While for the last group, the order is from top to bottom and from left to right.

(a) Original image

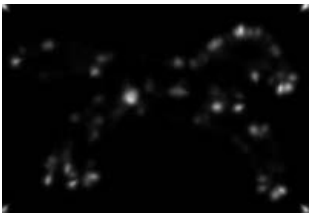

(b) First curvature matrix

(c) Second curvature matrix

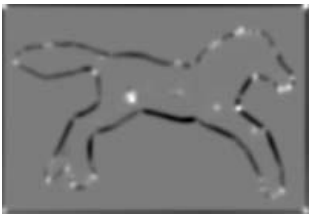

(d) Damped response from R function

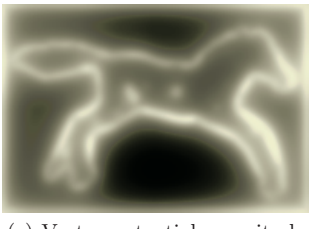

(e) Vector potential magnitude

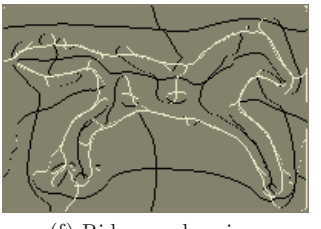

(f) Ridges and ravines

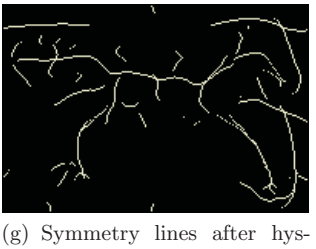

(g) Symmetry lines after hysteresis

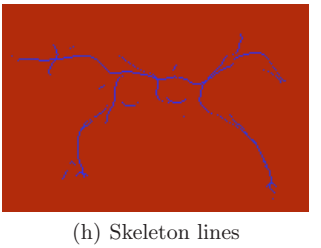

(h) Skeleton lines

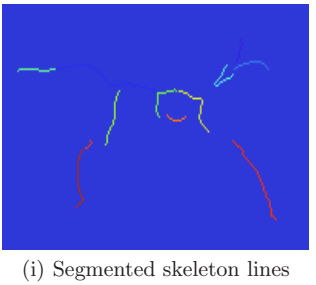

(i) Segmented skeleton lines

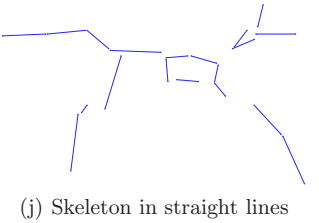

(j) Skeleton in straight lines

Fig. 1. Skeletonisation and noise elimination

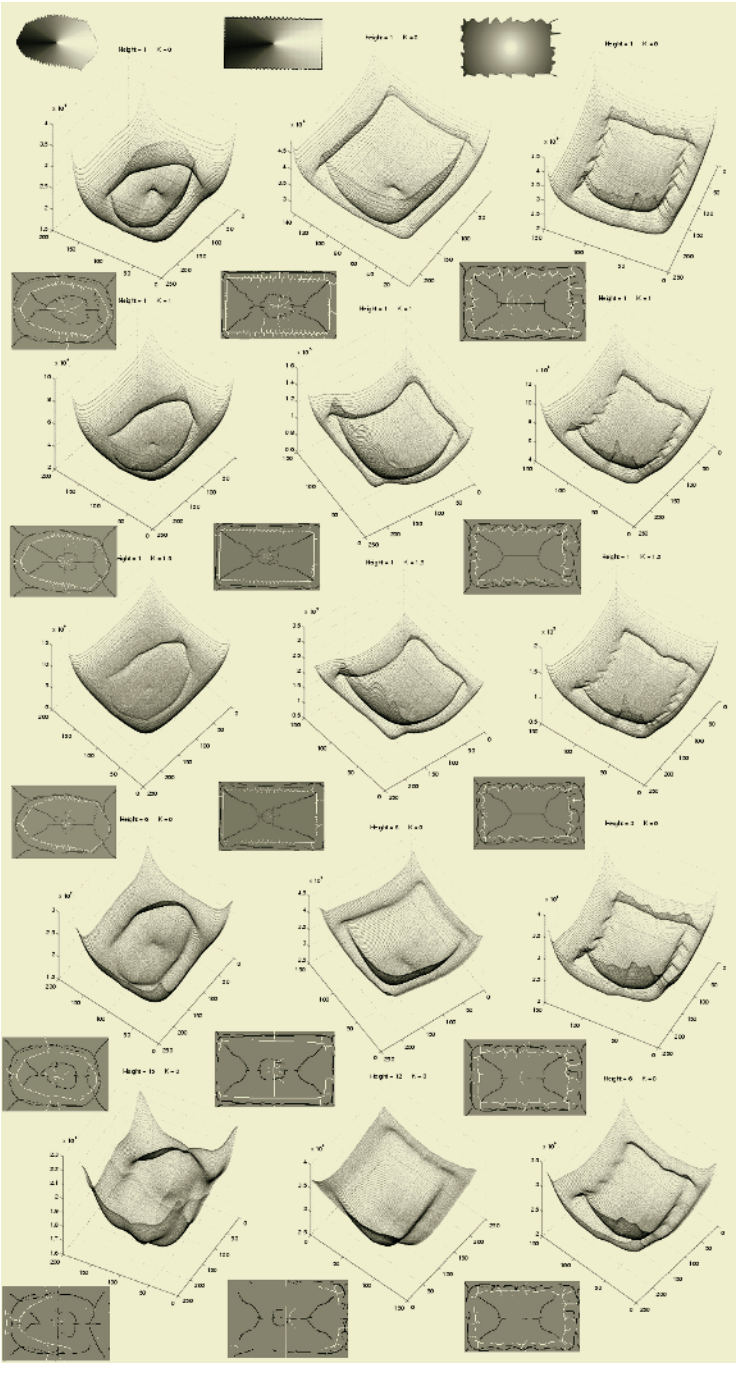

Fig. 2. Synthetic images with increasing K and height

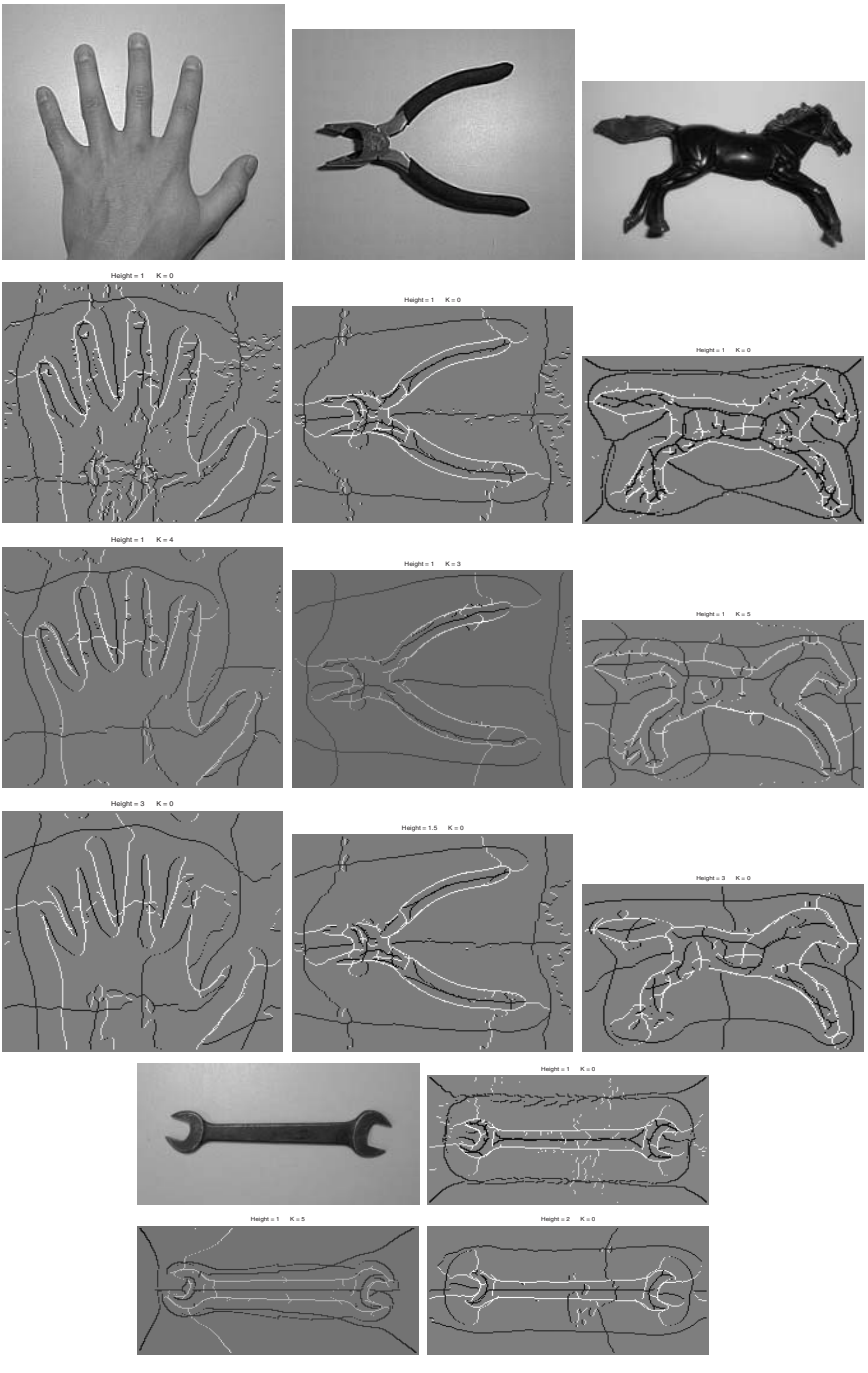

Fig. 3. Experimental results for real world objects

The main feature to note from these examples is that we can damp the noise by increasing parameter K as well as by increasing the height, but curvature damping outperforms height sampling in the following ways. First, the visual appearance of the results is more pleasing. Second, we preserve local curvature information. Third, as is the case with the synthetic images, when we increase the height, some useful ridges and ravines, especially ridges, begin to vanish rapidly as well as the noise.

6 Conclusion

In this paper, we incorporate curvature effects into the computation of grey-scale skeletons. The method builds on the magnetostatic analogy of Cross and Hancock, and employs a curvature dependent current damping. In addition, we have described a number of postprocessing steps that can be used to improve the quality of the detected skeletons. The advantages of the method are improved noise resilience of the detected skeleton, and better skeleton connectivity. In the end, we illustrate how to control high boundary noise by varying constant K rather than by increasing height. This results in good noise control and does not blur the details of the skeleton.

References

1. C. Arcelli and G.S Di Baja. A width-independent fast thinning algorithm. *IEEE PAMI*, 7(4):463–474, 1985.
2. H Blum. *A transformation for extracting new descriptors of shape.* MIT Press, W. Wathen-Dunn Ed. Models for the perception of speech and visual form, 1967.
3. A.D.J. Cross and E.R. Hancock. Scale space vector fields for symmetry detection. *Image and Vision Computing*, 17:337–345, 1999.
4. C.G. Harris and M.J. Stephens. A combined corner and edge detector. *Proceedings Fourth Alvey Vision Conference*, pages 147–151, 1988.
5. B. Luo, A. D. Cross, and E. R. Hancock. Corner detection via topographic analysis of vector potential. *Pattern Recognition Letters*, 20(6):635–650, 1999.
6. H. Moravec. Obstacle avoidance and navigation in the real world by a seeing robot rover. doctoral dissertation, Robotics Institute, Carnegie Mellon University, 1980.
7. R.L. Ogniewicz and O. Kübler. Hierarchic voronoi skeletons. *Pattern Recognition*, 28(3):343–359, 1995.
8. K. Siddiqi, S. Bouix, A. Tannenbaum, and S.W Zucker. The hamilton-jacobi skeleton. *International Conference on Computer Vision*, pages 828 – 834, 1999.
9. K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999.
10. Z. S. G. Tari, J. Shah, and H. Pien. Extraction of shape skeletons from grayscale images. *Computer Vision and Image Understanding*, 66:133–146, 1997.
11. H. Tek, P. A. Stoll, and B. B. Kimia. Shocks from images: Propagation of orientation elements. *IEEE CVPR*, pages 839–845, 1997.
12. Andrea Torsello and Edwin R. Hancock. A skeletal measure of 2d shape similarity. *Lecture Notes in Computer Science*, 2059:260–271, 2001.