Automatic image segmentation using a deformable model based on charged particles

Andrei C. Jalba, Michael H.F. Wilkinson, Jos B. T. M. Roerdink

Institute of Mathematics and Computing Science
University of Groningen, P.O. Box 800
9700 AV, Groningen, The Netherlands
{andrei,michael,roe}@cs.rug.nl
http://www.cs.rug.nl

Abstract. We present a method for automatic segmentation of grey-scale images, based on a recently introduced deformable model, the charged-particle model (CPM). The model is inspired by classical electrodynamics and is based on a simulation of charged particles moving in an electrostatic field. The charges are attracted towards the contours of the objects of interest by an electrostatic field, whose sources are computed based on the gradient-magnitude image. Unlike the case of active contours, extensive user interaction in the initialization phase is not mandatory, and segmentation can be performed automatically. To demonstrate the reliability of the model, we conducted experiments on a large database of microscopic images of diatom shells. Since the shells are highly textured, a post-processing step is necessary in order to extract only their outlines.

1 Introduction

An important aspect in many image analysis and computer vision tasks is image segmentation, the process in which an image is divided in its constituent parts. Here, we shall focus on boundary-based segmentation using the recently introduced charged-particle model (CPM) [1].

The CPM is inspired by classical electrodynamics and consists of a system of charged particles moving in an electrostatic field. The charges are attracted towards the contours of the objects of interest by an electric field, whose sources are computed based on the gradient-magnitude image. The electric field plays the same role as the potential force (defined to be the negative gradient of some potential function) in the snake model, while internal interactions are modeled by repulsive electrostatic forces (referred to as Coulomb forces). The method needs an initialization step, which is much less critical than in the snake model. Unlike the active contour model, in our model charges can be placed entirely inside an object, outside on one side of the object, or they can cross over parts of boundaries. In contrast to attractive forces based on the squared gradient-magnitude image [2], which act only in small vicinities along boundaries of objects, the electric field exhibits increased capture range because of its long range attraction, and enhanced robustness of the model against boundary leakage. Due to the combined effect of external interactions of particles with the electrostatic field, and internal repelling forces between them, particles follow paths along object boundaries, and hence

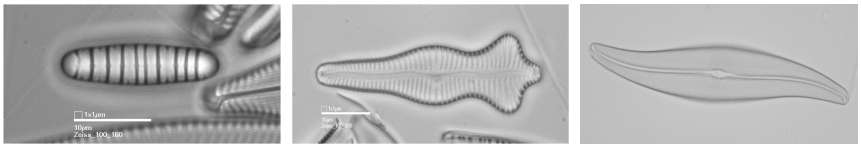

Fig. 1. Some examples of diatom shells.

converge without difficulty into deep boundary concavities or internal boundaries separating embedded objects. Moreover, the method is insensitive to initialization, and can adapt to topological changes of the underlying shape, see [1].

In this paper we present methods for automatic segmentation based on the CPM, using different strategies for automatic initialization of particles: (i) particles are spread uniformly over the image plane, (ii) particles are placed at locations of high gradient-magnitude, and (iii) particles are initialized on boundaries of the regions found by a marker-selection procedure [3]. To demonstrate the reliability of the model, we conducted experiments on a large database of microscopic images of diatom shells (see Fig. 1 for some examples).

2 The Charged-particle model (CPM)

The CPM consists of a system of N positively charged particles p_i with electric charges q_i, $i = 1 \dots N$, which freely move in an electrostatic field $\boldsymbol{E}$, generated by fixed, negative charges, placed at each pixel position of the input image, with charge magnitude proportional to the edge-map of the input image. Therefore, each free particle q_i moves under the influence of two forces: (i) internal Coulomb force, $\boldsymbol{F}_c$, due to the interaction of the particle with other free particles, and (ii) external Lorentz force, $\boldsymbol{F}_l$, due to the electric field generated by the fixed negative charges e_i, see Fig. 2. The resulting force $\boldsymbol{F}$ acting on a particle p_i located at position vector $\boldsymbol{r}_i = [x_i, y_i]$ is

$$\boldsymbol{F}(\boldsymbol{r}_i) = \boldsymbol{F}_c(\boldsymbol{r}_i) + \boldsymbol{F}_l(\boldsymbol{r}_i), \tag{1}$$

where $\boldsymbol{F}_c$ is the Coulomb force and $\boldsymbol{F}_l$ is the Lorentz force. Assuming that all free particles have the same positive charge $q_i = q$, it can be shown that the equilibrium equation (Eq. (1)) can be rewritten as

$$\boldsymbol{F}(\boldsymbol{r}_i) = w_1 \sum_{j \neq i}^{N} \frac{\boldsymbol{r}_i - \boldsymbol{r}_j}{|\boldsymbol{r}_i - \boldsymbol{r}_j|^3} - w_2 \sum_{k:\boldsymbol{R}_k \neq \boldsymbol{r}_i}^{M} e_k \frac{\boldsymbol{r}_i - \boldsymbol{R}_k}{|\boldsymbol{r}_i - \boldsymbol{R}_k|^3}, \tag{2}$$

where $w_1 = kq^2$ and $w_2 = kq$ are weights, and k is a constant. The major difference between the two terms in Eq. (2) is that the Lorentz force reflects particle-mesh or external interactions and is computed in the image domain, while the Coulomb force represents particle-particle or internal interactions. Therefore, each particle is the subject of two antagonistic forces: (i) the Coulomb force, which makes the particles to repel each other, and (ii) the external Lorentz force which attracts the particles. Since

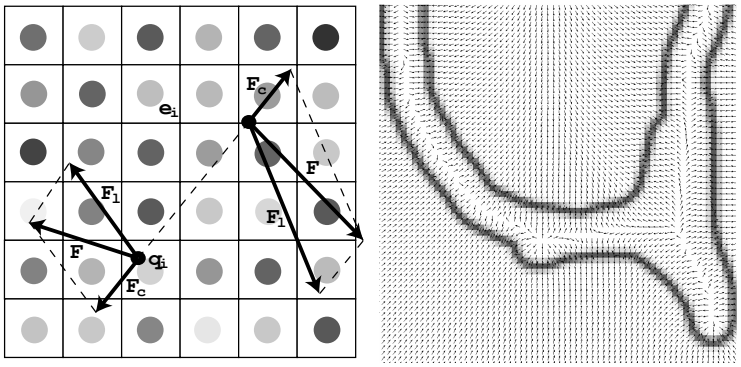

Fig. 2. The charged-particle model. *Left*: Forces acting on free particles q_i (indicated by small black dots) which move in the electric field generated by fixed charges e_i (indicated by grey dots); different grey values represent different charge magnitudes. *Right*: Example of electrostatic field $\boldsymbol{E}$ generated by fixed charges.

the distribution of fixed charges e_i reflects the strength of the edge map, and the electric force is "inverse-square", i.e., it decays with the squared distance, the electrostatic field has large values near edges and small values in homogeneous regions of the objects present in the input image.

2.1 Particle dynamics

The total energy of the system is the summation of all particle energies, i.e.,

$$E_p(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_N) = \\ \frac{1}{2}\sum_{i=1}^{N}\left(w_1 \sum_{j\neq i}^{N} \frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} - w_2 \sum_{k:\boldsymbol{R}_k \neq \boldsymbol{r}_i}^{M} \frac{e_k}{|\boldsymbol{r}_i - \boldsymbol{R}_k|} \right). \tag{3}$$

Having defined the energy associated with our system, we can derive its equations of motion. The variations of particle potentials with respect to positions produce forces acting on particle positions. The standard approach is to consider the Newtonian equations of motion, and to integrate the corresponding system of differential equations in time, i.e.,

$$\boldsymbol{F}(\boldsymbol{r}_i) = w_1 \boldsymbol{F}_c(\boldsymbol{r}_i) + w_2 \boldsymbol{F}_l(\boldsymbol{r}_i) - \beta \boldsymbol{v}_i \tag{4}$$

$$\boldsymbol{a}_i = \frac{\boldsymbol{F}(\boldsymbol{r}_i)}{m_i} = \frac{d^2 \boldsymbol{r}_i(t)}{dt^2}, \tag{5}$$

where m_i is the mass of the particle p_i (we set $m_i = 1$), and $\boldsymbol{r}_i$, $\boldsymbol{v}_i$ and $\boldsymbol{a}_i$ are its position, speed and acceleration, respectively. Notice that compared to Eq. (1), Eq. (4) has an additional term, $\boldsymbol{F}_{damp}(\boldsymbol{r}_i) = -\beta \boldsymbol{v}_i$, the damping (or viscous) force which is required by the particles to attain a stable equilibrium state, which minimizes their

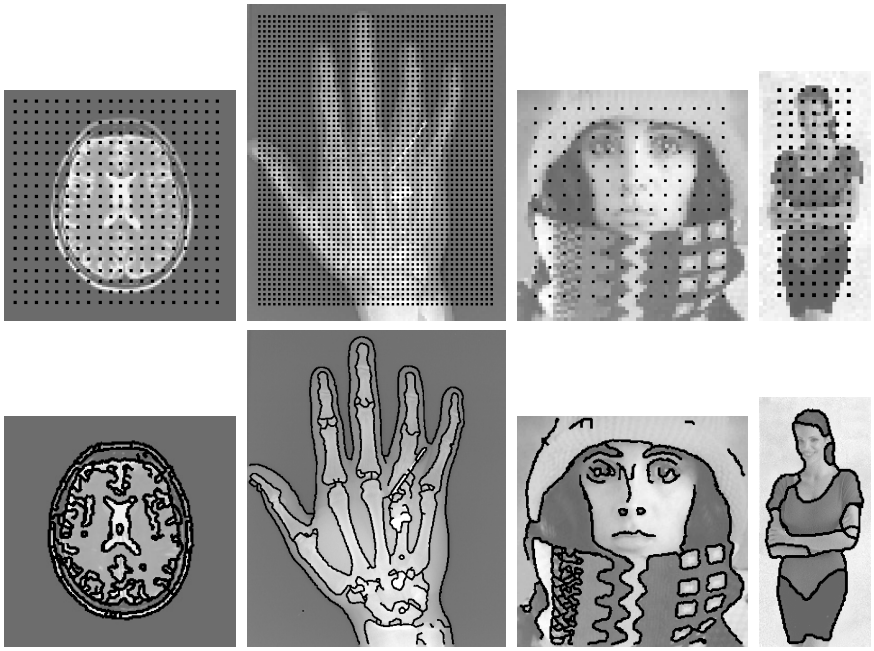

Fig. 3. Automatic segmentation. *First row*: initializations; *second row*: segmentation results.

potential energies, see Eq. (3). Eq. (5) is written as a system of coupled, first order differential equations, and solved using some method for numerical integration [1, 4]. For detailed information on the CPM, efficient methods for implement it and pseudo-code, we refer to [1].

2.2 Curve reconstruction

So far, our particle system does not provide us with explicit representations of object boundaries. This problem can be thought of as that of curve reconstruction from unorganized points: we are given a set of points and asked to connect them into the most likely polygonal curve.

If the aim is to recover only one, closed contour, the reconstruction problem can be formulated as enumerating the particles and then ordering them into a sequence which describes a closed contour along the boundary of the object. The problem is now isomorphic to the classical symmetric traveling salesman problem (STSP), and established techniques for approximating TSP can be used. This approach can be useful in assisted (interactive) segmentation of medical imagery, where the very purpose of segmentation can be to isolate and extract a specific object (e.g. a tumor). However, under the more general assumption that no a priori knowledge about the underlying topology is available, curve reconstruction algorithms must be involved. Therefore, in all experiments reported below, we use the algorithms by Amenta et al. [5] to reconstruct the recovered curves.

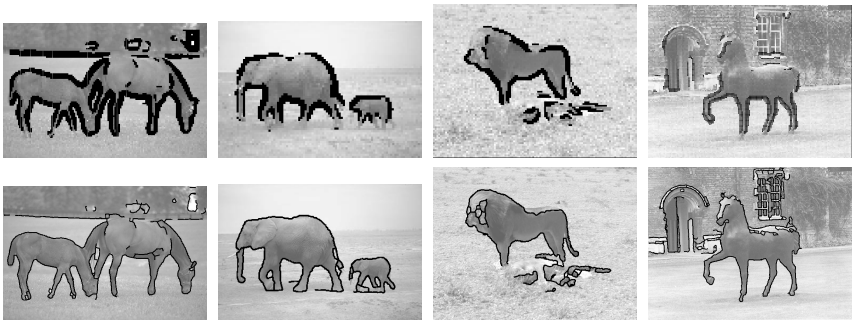

Fig. 4. Segmentation of natural images. *First row*: initializations; *second row*: results.

3 Segmentation results

3.1 Natural images

Our first experiment is automatic segmentation of natural images using (trivial) automatic strategies for initialization. In all experiments which we report in this paper, we used the same values for the two weights w_1 and w_2 (see Eq. (2)), i.e. $w_1 = 0.6$ and $w_2 = 0.7$, and all other parameters of the model were set as in [1]. The pre-processing step consists in image filtering by means of a Gaussian pyramid with three levels.

With this experimental setup, the first set of segmentation results is shown in Fig. 3. The initializations shown in this figure were performed by uniformly spreading particles over the image plane. As it can be seen, the most important structures present in these images were correctly recovered.

The second set of results is shown in Fig. 4. In this case, free particles were placed at those locations of the gradient-magnitude image with values above 10% of the maximum magnitude. Natural images are known to be particularly difficult to segment, mostly because of the background texture surrounding the main objects. Without being perfect, the segmentation results shown in both figures are quite good, even though a very simple initialization method was used.

3.2 Results for a large database of diatom-shell images

The second experiment we conduct is automatic segmentation, on a large database consisting of 808 diatom images (see Fig. 1 for some examples).

The goal is to extract the outline of each diatom shell present in the input image. The extracted outlines, encoded as chain-codes , provide the input for identification methods such as those in [6].

The input consists of grey-scale, high-magnification images of diatom shells obtained by automatic slide scanning [7]. Ideally, each image contains a single diatom shell, but as it can be seen in the figure, diatoms may lay on top of each other, may not be in proper focus, or they can be very close to each other. Moreover, dust specks and background texture may be visible in some images.

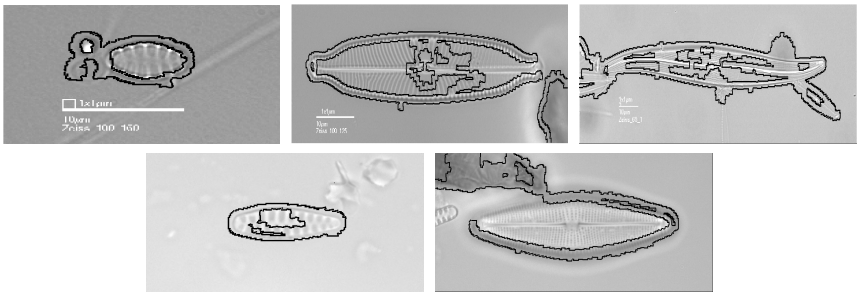

Fig. 5. Problematic diatom images for the CPM, with superimposed initializations.

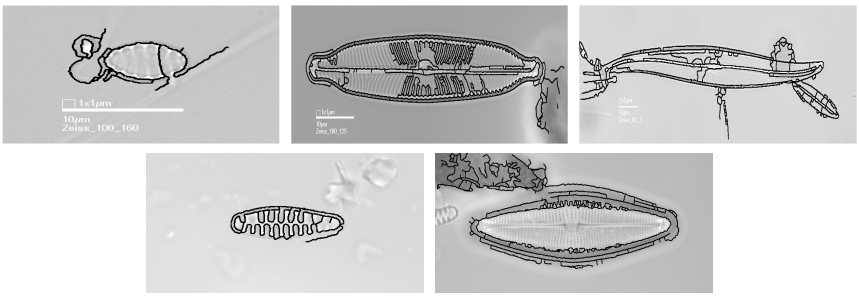

Fig. 6. Problematic diatom images for the CPM; final (erroneous) results.

Most diatoms in images such as those in Fig. 1 present prominent outlines which can be detected either by thresholding or by edge detectors. Unfortunately, if the illumination around the diatom is not uniform, most global thresholding methods fail to find a proper threshold value. In addition, in microscopic images, diatoms exhibit the same grey levels as the background, and the histogram is unimodal [8]. This fact upsets most threshold selection methods which make the assumption that the histogram of the image is multimodal. Moreover, if the diatom is not in proper focus, the edges are blurred, and can only be partly detected by most edge detection techniques. Therefore, we use a method based on morphological filtering [3] to provide marker-regions (the same method was used in [3] in the context of watershed-based segmentation), and we initialize the particles on the boundaries of these regions. To guarantee that only one closed contour per diatom is extracted, each contour obtained using a standard contour-following algorithm is flood-filled, and then, traced once again.

With this experimental setup, the method succeeded in extracting 99.4% of visually-estimated correct contours. The initializations and final results (without the contour-tracing step) for the five cases in which the method failed are shown in Figs. 5 and 6, respectively. Four of the images shown in Fig. 5 have debris or fragments of other diatoms very close to the central diatom. The fourth image shows a very low contrast of the diatom outline, which is reflected in the weak gradient-magnitude response that is used by the CPM to compute the electric field. Nevertheless, in our opinion this is a very good result, considering that the CPM is a boundary-based method.

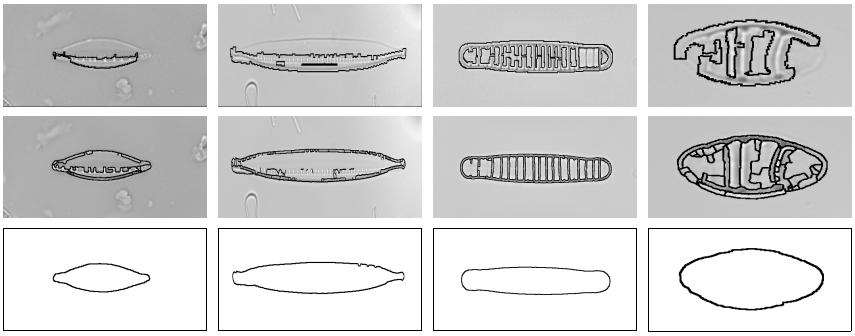

Fig. 7. Difficult diatom images (and initializations), correctly segmented by the CPM. *First row*: initializations; *second row*: reconstructed curve(s); *third row*: extracted diatom contours.

Fig. 7 shows some example results obtained using the CPM on difficult images on which a hybrid technique based on the morphological watershed from markers failed. This method obtained 98% (i.e. 16 errors), of correctly extracted contours, see [3].

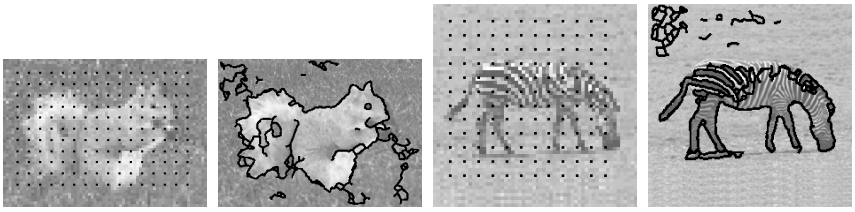

Fig. 8. The CPM may fail if highly textured regions surround the main object or belong to the main object; results obtained with the first initialization method.

3.3 Discussion

The advantages of using the second and third initialization strategies over the first one are twofold. First, the particles are already close to the final equilibrium positions, and therefore the total convergence time is smaller. Second, using the first initialization method, it may happen that some particles will be attracted towards highly textured regions, which are also regions with high response of the gradient magnitude, and therefore they will be trapped at these regions, see Fig. 8.

Fig. 9 shows segmentation results obtained using the second and third initialization strategies; see also the result in Fig. 3 obtained with the first method. The CPU timings (on a Pentium III machine at 670 MHz) for segmenting this x-ray image of 417×510 pixels were 45, 20, 25 seconds, using the first, second and third initialization methods, respectively.

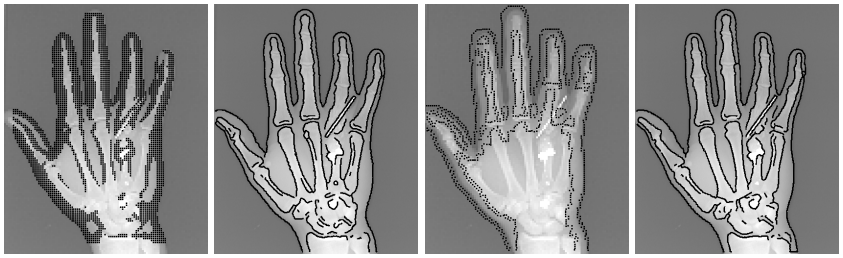

(a) (b) (c) (d)

Fig. 9. Comparative segmentation results; (a) initialization by the second method, (b) result, (c) initialization using the third method, (d) result.

4 Conclusions

The experimental results presented in this paper showed that the CPM can be used successfully to perform automatic segmentation, provided that a suitable setup has been identified.

Further investigations of the CPM are the subject of ongoing research. We shall focus on supplementing the energy formulation of the model with some information useful in the reconstruction phase. A shortcoming of the current method is that it cannot guarantee that the recovered contours (surfaces) are without gaps. Finally, many improvements of the CPM are possible. For example, instead of using Gaussian pyramids, one can use wavelet or other pyramids based on non-linear diffusion operators.

References

1. Jalba, A.C., Wilkinson, M.H.F., Roerdink, J.B.T.M.: CPM: A deformable model for shape recovery and segmentation based on charged particles. IEEE Trans. Pattern Anal. Machine Intell. (2004) in press.
2. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. Int. J. Comput. Vis. **1** (1987) 321–331
3. Jalba, A.C., Roerdink, J.B.T.M.: Automatic segmentation of diatom images. In: Proc. Comput. Anal. Images Patterns 2003. Volume 2756 of Lecture Notes in Computer Science. (2003) 369–376
4. Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T.: Numerical Recipes in C: The Art of Scientific Computing. Cambridge Univ. Press, Cambridge (1988)
5. Amenta, N., Bern, M., Eppstein, D.: The crust and the β-skeleton: Combinatorial curve reconstruction. Graphical Models and Image Processing **60** (1998) 125–135
6. Wilkinson, M.H.F., Jalba, A.C., Urbach, E.R., Roerdink, J.B.T.M.: Identification by mathematical morphology. In Du Buf, J.M.H., Bayer, M.M., eds.: Automatic Diatom Identification. Volume 51 of Series in Machine Perception and Artificial Intelligence. World Scientific Publishing Co. , Singapore (2002) 221–244
7. Pech-Pacheco, J.L., Cristobal, G.: Automatic slide scanning. In du Buf, H., Bayer, M.M., eds.: Automatic Diatom Identification. World Scientific Publishing, Singapore (2002) 259–288
8. Fischer, S., Bunke, H., Shahbazkia, H.R.: Contour extraction. In du Buf, H., Bayer, M., eds.: Automatic Diatom Identification. World Scientific Publishing, Singapore (2002) 93–107