

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Adaptive shape prior for recognition and variational segmentation of degraded historical characters

Itay Bar-Yosef^{a,*}, Alik Mokeichev^a, Klara Kedem^a, Itshak Dinstein^b, Uri Ehrlich^c

^aDepartment of Computer Science, Ben-Gurion University, Israel

^bDepartment of Electrical and Computer Engineering, Ben-Gurion University, Israel

^cDepartment of Jewish thought, Ben-Gurion University, Israel

ARTICLE INFO

Article history:

Received 16 August 2008

Accepted 15 October 2008

Keywords:

Segmentation

Degraded character recognition

Historical documents

Level set

Shape prior

ABSTRACT

We propose a variational method for model based segmentation of gray-scale images of highly degraded historical documents. Given a training set of characters (of a certain letter), we construct a small set of shape models that cover most of the training set's shape variance. For each gray-scale image of a respective degraded character, we construct a custom made shape prior using those fragments of the shape models that best fit the character's boundary. Therefore, we are not limited to any particular shape in the shape model set. In addition, we demonstrate the application of our shape prior to degraded character recognition. Experiments show that our method achieves very accurate results both in segmentation of highly degraded characters and both in recognition. When compared with manual segmentation, the average distance between the boundaries of respective segmented characters was 0.8 pixels (the average size of the characters was $70 * 70$ pixels).

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

Much effort has been devoted in the last few years to the digitization of books and documents in order to archive them in digital form. A common need in libraries and archives is to improve the readability of historical documents. This has high cultural and scientific values, e.g., enhancing degraded documents, improving OCR accuracy and facilitating paleographic researchers for which the accuracy of the segmentation process is of high research value. However, segmentation of historical documents is difficult due to varying contrast, smudges, faded ink, and the presence of bleed-through text. In order to overcome these difficulties, binarization algorithms designed especially for historical documents have been proposed in the last few years [1,2]. Although for some cases impressive results were obtained, these methods often fail when dealing with extremely degraded documents. The main reason is that most of the binarization methods are based on global or local statistics derived from the gray-scale image. These statistics alone are not sufficient in complex cases, where characters are broken and/or partially erased as shown in Fig. 1.

A small number of scientific papers have dealt with restoration of degraded historical documents. Very few among them considered gray-scale images. Droettboom [3] and Hobby [4] dealt with binary documents. Droettboom et al. [3] proposed a method based on graph combinatorics to merge broken characters of historical documents. The goal of their algorithm was to find an optimal way to join connected components on a given page, that maximizes the mean confidence of all characters. Hobby et al. [4] proposed to improve the quality of degraded text images by image matching techniques. Similar symbols were clustered, and a prototype of each cluster was generated to replace the cluster symbols. An active contour model for restoration of degraded gray-scale characters was developed by Allier et al. [5]. The authors incorporated a single shape model to a parametric active contour based on a gradient vector flow (GVF) representation [6] of the image and of the shape prior. Although impressive results have been reported, their approach does not provide satisfying results when applied on highly degraded text image with large variability of the characters shapes.

In this paper, we present a variational approach for accurate segmentation of highly degraded characters. Our main contribution is a novel adaptive shape prior that is customized to the character's gray-scale image and is not limited to any particular shape of the training set. We integrate the shape prior into a variational segmentation model and also demonstrate how the proposed shape prior can significantly improve degraded character recognition.

* Corresponding author.

E-mail addresses: itaybar@cs.bgu.ac.il (I. Bar-Yosef), mokeiche@cs.bgu.ac.il (A. Mokeichev), klara@cs.bgu.ac.il (K. Kedem), dinstein@ee.bgu.ac.il (I. Dinstein), ehrllich@exchange.bgu.ac.il (U. Ehrlich).



Fig. 1. Example of highly degraded text.

The rest of the paper is organized as follows: In Section 2 we present our shape prior and describe the prior construction process. Section 3 demonstrate the application of our shape prior with respect to recognition. Section 4 gives a short introduction on active contours and level set methods, and outlines the variational formulation of the shape prior segmentation. Experimental results are presented in Section 5, and in Section 6 we conclude our work and outline the future work.

2. Definition and construction of the shape prior

Accurate binarization of highly degraded characters is often practically impossible without using any prior knowledge of the character's shape or identity. A simple method to obtain the degraded character's identity is matching based on normalized cross-correlation (NCC) which has been proven to be robust against vast gray-scale degradations. The main drawbacks of the NCC is its weakness against local geometric distortions and affine-transformations [7]. However, as commonly seen in historical documents, the main differences in character shapes are due to the writer's writing style, thus the main challenge is to overcome these local differences while affine-distortions (orientation and scale) are minor and can be neglected. We note that affine-invariance can be achieved by using more sophisticated approaches such as the one proposed by Wakahara et al. [7].

In the following, we propose a novel shape model which is adapted to the character's boundaries very accurately. The shape prior construction process described below assumes a given gray-scale character image and a small set of shape models corresponding to the identity of that character. In Section 3 we relax the assumption of known character's identity and show how to utilize the shape prior for character recognition. The creation of the shape models used for prior construction is further discussed in Section 2.3.

Given a gray-scale image of a certain character and a set of N respective shape models, we align each shape to the gray-scale image based on NCC (each model is aligned according to their maximum correlation lags). Naturally, each model fits the gray-scale image differently (see example in Fig. 2). We exploit this fact to create a new shape model constructed of the models' fragments that best fit to the character's boundary. The main advantage of our approach is that our shape model captures local information of the gray-scale character image with respect to the given models. We further use this shape model for both accurate segmentation and recognition of highly degraded characters (see Sections 3 and 4).

2.1. Confidence map construction

Given a gray-scale image and a set of N respective shape models, we define a mechanism that provides local information on the degree of fitting of each model to the gray-scale image. We will further utilize this information for local comparison between the models in order to find those fragments that best fit to the image.

For a given gray-scale image and an aligned shape model, we create a *confidence map* in two steps. First, we superimpose the model on the gray-scale image. Around each point along model's boundary, we place a small window (Fig. 3(a)) and calculate the normalized correlation between the corresponding portion of the model and that of the image (Fig. 3(b)). The calculated value is then assigned to the corresponding contour point (Fig. 3(c)). In order to compare between the different models locally, we propagate the score from each model's boundary to the rest of the image domain. For that purpose, we apply the fast image inpainting algorithm proposed by Oliveira et al. [8]. Their algorithm iteratively updates unknown pixel values with the average value of their known surrounding neighbors. Propagating the fitting score from the boundary, results in a confidence map which measures at each pixel, the local fitting to the model. For the set of N shape models, we define s_i to be the *confidence map* of the i th model. Fig. 3(d) presents the confidence map of the shape model shown in white in Fig. 3(a). In the following section, we describe our shape representation model and explain how we use the confidence maps of the different models to construct the customized shape prior.

2.2. Shape representation

After calculating the confidence map of the N respective shape models, we determine for each pixel in the image the most suitable model, i.e., the one with the highest confidence value. The shape prior is composed from pixels of the models such that each model contributes only pixels that best fit the raw data.

For this purpose we use the signed distance function (SDF) for shape representation. This representation has gained much popularity in the recent years, mainly in the field of variational segmentation. We will elaborate on the advantages of this representation in Section 4. The idea is to represent the shape's contour C by embedding it in a higher dimension level set function Φ , as follows:

$$\Phi(x) = \begin{cases} -D(x, C), & x \text{ is inside } C \\ D(x, C), & x \text{ is outside } C \\ 0, & x \in C \end{cases} \quad (1)$$

where $D(x, C)$ denotes the Euclidean distance from x to the closet point on C . The contour C can be reconstructed from such representation by taking its zero level set $C = \{x | \Phi(x) = 0\}$. An example for SDF is shown in Fig. 4(b). Denote by Φ_i the SDF of the i th model, and denote by Φ_p the composite shape prior for the character. For each pixel (x, y) in the image, the model with the highest confidence value determines Φ_p :

$$\Phi_p(x, y) = \Phi_k(x, y), \quad k = \underset{i}{\operatorname{argmax}}(s_i(x, y)) \quad (2)$$

In a similar way to Φ_p , we define the *prior confidence map* η , which measures the reliability of the composite shape prior at each pixel:

$$\eta(x, y) = s_k(x, y), \quad k = \underset{i}{\operatorname{argmax}}(s_i(x, y)) \quad (3)$$

Fig. 5 shows the shape prior Φ_p and the confidence map $\eta(x, y)$ based on the gray-scale image and the three shape models shown in Fig. 2. Notice the high correspondence between the shape prior fitting to the gray-scale image, and the confidence map. In Section 4

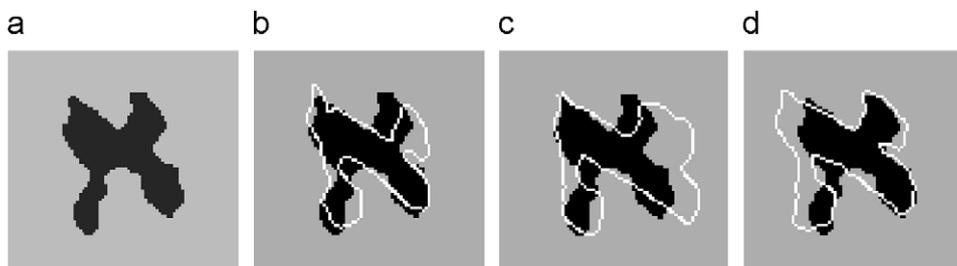


Fig. 2. (a) Gray-scale image of the Hebrew letter "Aleph". (b–d) Different models (superimposed in white) aligned based on the maximum normalized cross-correlation. Although none of these models has an exact match with the gray-scale character's boundary, by taking from each model only the fragments that best fit the boundary locally, an accurate approximation to the original character's shape could be achieved.

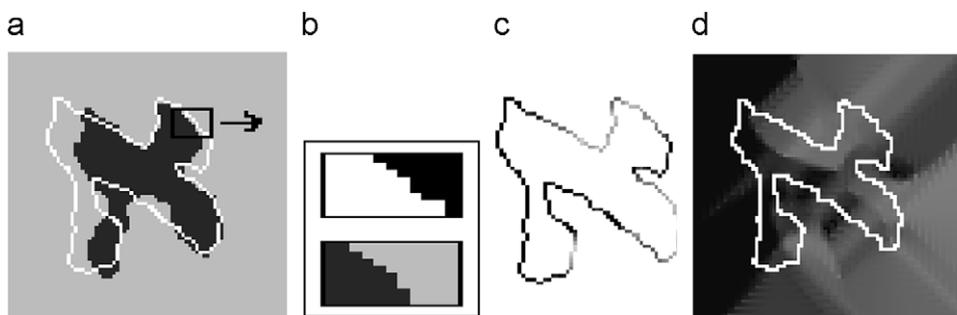


Fig. 3. Confidence map creation. (a) A small rectangular window (in black) is placed on each boundary point of the model. (b) Normalized correlation is then calculated on the extracted image and model region. (c) Scores on the model's boundary. (d) The resulting confidence map s_i , after propagating the information from the model's boundary (the model's boundary is superimposed in white). Both in (c) and (d) bright regions indicate high fitting score, while dark regions indicates poor fitting.

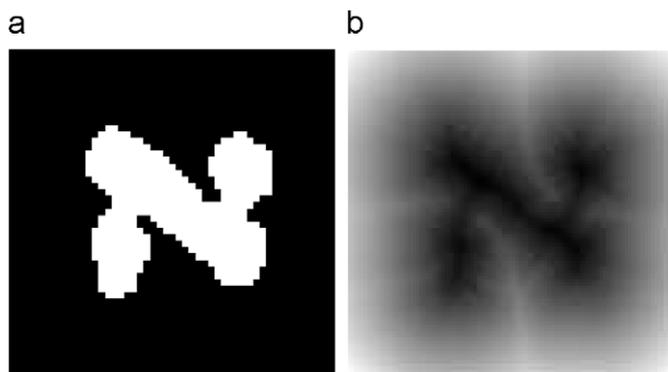


Fig. 4. (a) Binary shape image. (b) Corresponding SDF Φ . The SDF image is shown in gray-scale values, where negative values (inside the shape) are shown in dark colors and positive values (outside the shape) are shown in bright ones.

we will use this information to provide accurate segmentation of the character.

2.3. Model set creation

Given a training set of n characters of the same letter, we construct a smaller set of N shape models from which the adaptive shape prior to be constructed. In principle, the entire sample of n characters could be used as models and the reduction of set size is only for computational efficiency. Ideally, the smaller set of models should have minimal redundancy and should closely represent the variations that exist among the shapes of characters. For this purpose we identify clusters of characters having similar shapes and use the averaged shape of each of the clusters as a model. The shape models

are constructed as follows:

1. Compute the SDF of the training characters, and rearranged each SDF as a column vector.
2. Apply principal component analysis (PCA) to reduce the SDF vectors dimensionality. In our experiments, SDF vectors of 4900 components ($70 * 70$ images) were represented with 95% accuracy by 10 principal components.
3. Cluster the n 10-dimensional vectors into N clusters.
4. For each cluster, obtain a new SDF by averaging the SDF's of characters belonging to that cluster.
5. Use the zero level sets of the N obtained SDFs as shape models.

Fig. 6(a) shows a set of 20 characters from which we construct four models through clustering in the reduced dimension character space. Fig. 6(b) shows the constructed models. The choice of four models is rather arbitrary, and using a larger number of models for prior construction is expected to improve the accuracy of representations. It turns out that four models are sufficient for obtaining highly accurate reconstructions as we demonstrate in Section 5.

3. Application of the proposed shape prior to recognition

One of the simplest forms of character recognition is based on NCC. In this approach, a prototype letter (usually the average shape) is chosen for each letter class and the prototype having the highest correlation with the raw gray-scale image determines character's identity. As mentioned before, one of the main drawbacks of the NCC is its sensitivity to local distortions, or in our case, the variability of the writer's writing style. In our data, additional sources of difficulty exist, such as stains, erased letter parts and highly variable contrast within a single character. These indeed significantly decrease the accuracy of the recognition. An experiment conducted on 500 letters of 16 different letter classes yielded only 77% correct recognition.

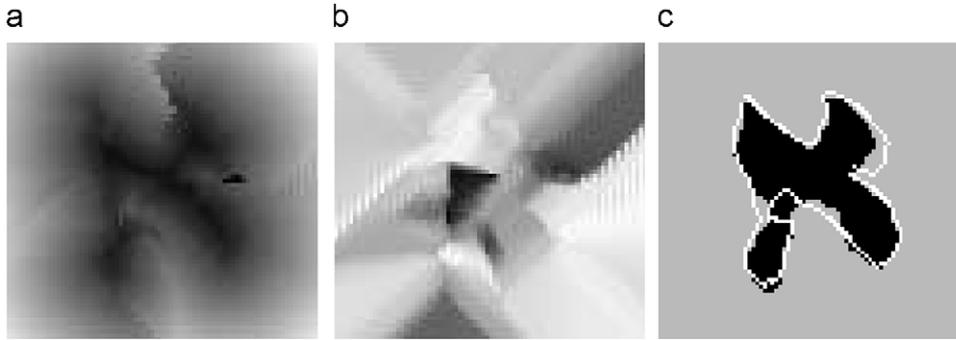


Fig. 5. (a) The shape prior Φ_p , calculated based on the gray-scale image and the three models in Fig. 2. (b) The prior's confidence map η . (c) The zero level set of Φ_p (in white) super imposed on the gray-scale image. In regions where the shape prior fits accurately, the confidence map has high values. Poor fitting of the models to the image results in low confidence values.

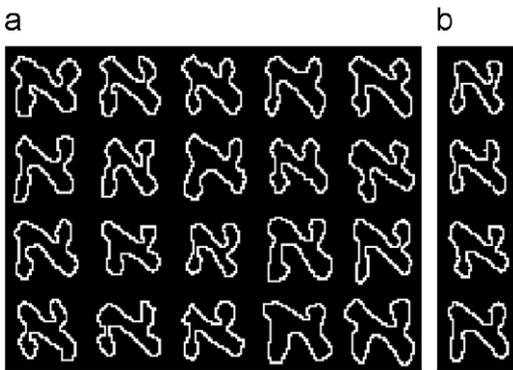


Fig. 6. (a) A training set of 20 characters from which we construct a compact set of four models. (b) The constructed models.

Our prediction is that results should significantly improve if recognition is based on the customized shape priors rather than on predetermined fixed templates. Given a gray-scale image for recognition, a set of shape priors is created, one for each letter class, as described in the previous section. Our underlying assumption is that the best matching shape prior (highest correlation with the gray-scale image) would be created from the models belonging to the true class of the character to be recognized.

Analyzing the results obtained with NCC matching using a predefined template, we noticed that the true character identity is rarely ranked after the best five matches. Therefore, we apply our shape prior composition process only on the best five matching identities.

Denote by Φ_i , $i = \{1, \dots, 5\}$ the composite shape prior of the first five matches. Since Φ_i is an SDF, we can extract its corresponding binary shape b_i , i.e., $b_i = \{x | \Phi_i(x) < 0\}$. We then calculate for each binary model its correlation with the underlying gray-scale image. The model having the highest correlation determines character's identity. This approach yielded a significant improvement over the same set of 500 characters with 98% correct recognition. We note that our approach can be easily adopted for more sophisticated recognition methods instead of the NCC. This can improve the recognition rate, and reduce the number of shape priors to be constructed (five in our case).

In the following section, we describe our approach for accurate variational segmentation based on our composite shape prior.

4. Segmentation with active contours and level set framework

In the last decade, variational methods have gained much popularity in the computer vision community and affected considerably

the approach to image segmentation. First introduced by Kass et al. [9], active contours can be divided into two main categories: edge based and region based models. Edge based models [10] are defined such that the active contour is attracted to high gradient image boundaries. On the other end, region based models [11,12] are designed such that the image is partitioned into several homogeneous regions. *Geodesic active contours*, a prominent edge based model, was proposed by Caselles et al. [10]. They showed that the original Kass's model can be formulated as finding a curve C with minimal geodesic length, i.e.

$$\min \int g(|\nabla I(C(s))|) ds \quad (4)$$

where g is an inverse edge indicator function (e.g., $g = 1/(1 + |\nabla I|^2)$). The steady state of the active contour can be reached by evolving each contour point according to the following flow:

$$\frac{\partial C}{\partial t} = g\kappa\bar{N} - (\nabla g \cdot \bar{N})\bar{N} \quad (5)$$

where κ is the contour curvature, and \bar{N} is the unit normal to the curve. Although the geodesic active contour has been widely used in different domains, it has two major drawbacks. First, the evolution is based on local information, therefore it is very sensitive to local minima. Second, the segmentation results are highly dependent on the initial position of the contour.

Region based models, in general, are more robust to textured and noisy data and are less sensitive to the initialization of the active contour. In Ref. [11], Chan and Vese proposed a region based model for variational segmentation based on the well known Mumford–Shah functional [13]. The main idea behind their approach is to partition the image into two separated homogeneous regions, defined by the evolving contour C (inside and outside the contour) which minimizes the following functional:

$$F(C, c_1, c_2) = \int_{C_{in}} |I - c_1|^2 dx dy + \int_{C_{out}} |I - c_2|^2 dx dy \quad (6)$$

where c_1 , c_2 are the mean intensity values inside and outside the contour C , respectively. In order to control the smoothness of their contour, Chan and Vese added the length of the contour $\int ds$, as a regularization term. In our implementation, we used the weighted length $\int g(C(s)) ds$, the *geodesic active contour* (Eq. (4)), which provides better segmentation results [14].

Both of the *geodesic* model and the *Chan–Vese* model implementations are based on the *level set* method, a numerical technique for front propagation proposed by Osher and Sethian [15]. The main idea of the level set method is to embed the evolving contour C as the zero level set of an implicit function Φ (usually the SDF), defined in

a higher dimension $C(t) = \{(x, y) | \Phi(t, x, y) = 0\}$. Then, the curve's evolution equation $\frac{\partial C}{\partial t} = F\vec{N}$ can be re-written in a level set formulation $\frac{\partial \Phi}{\partial t} = F|\nabla \Phi|$, where F is the evolution speed function. Topological changes of the evolving contour are handled naturally and efficient numerical techniques can be applied [15]. The level set formulation of Eqs. (5) and (6) are, respectively:

$$\frac{\partial \Phi}{\partial t} = \delta(\Phi) \operatorname{div} \left(g(|\nabla I|) \frac{\nabla \Phi}{|\nabla \Phi|} \right) \quad (7)$$

$$\frac{\partial \Phi}{\partial t} = \delta(\Phi) (-I - c_1)^2 + (I - c_2)^2 \quad (8)$$

where c_1, c_2 are updated with the level set evolution. $\delta(\cdot)$ is a smoothed Dirac function defined as in Ref. [16] and $\delta(\Phi)$ act as a narrow band around the zero level set.

Both the *geodesic* and the *Chan-Vese* models have been widely used in the last decade and established a solid ground for many variations of their algorithms. However, since both of these methods rely entirely on gray-scale intensities, they often fail when encountering noisy, degraded and occluded data. In such cases, *prior* information on the object to be segmented can significantly improve segmentation results.

The level set framework has been found to be very useful in integrating *prior* knowledge to the segmentation process, and many algorithms have been developed for this purpose. Chen et al. [16] proposed to use the average shape of a training set as a shape prior. A statistical approach is presented in the work of Leventon et al. [17], where the shapes of the training set are represented by their SDF and a shape prior is constructed from a reduced dimension space. Cremers et al. [18] proposed a variational model that incorporates a level set representation of the shape prior into Chan-Vese's algorithm [11]. The authors also introduced a labeling function which determines the region to which the shape prior should be applied. An improved version of Cremers work was proposed by Chan et al. [19]. For a recent review on incorporation of shape priors to active contours, please see Ref. [20].

Most of the prior based approaches treat the shape prior as a global term and are limited to a predefined set of permitted shapes. The main contribution of this paper is that we choose from each shape model only the fragments which fit the image best, and create a composite shape prior customized to the character's gray-scale image. Therefore, we are not limited to any particular shape of the training set.

In the following section, we describe how our composite shape prior described in Section 2 can be efficiently integrated in the level set segmentation model.

4.1. Incorporating the shape prior into the active contour

We formulate a variational segmentation model which is based on our composite shape prior Φ_p and the Chan-Vese model (Eq. (8)) described in Section 4. We use the *geodesic active contour* (Eq. (7)) as a regularization term to smooth the evolving contour.

Paragois et al. [21] represented a shape by its respective (SDF). The idea of embedding the shape prior in a level set function (SDF) has been shown to provide a powerful tool for integration in the level set framework [18,21]. Since the evolving contour is also represented as a level set function, the integration of the (SDF) based shape prior becomes natural.

The proposed shape prior Φ_p is build such that it defined an approximation to an SDF (Section 2.2). Therefore, we follow Paragois et al. [21] approach for the shape integration. Let Φ be the evolving level set function for the segmentation, and let Φ_p be our proposed

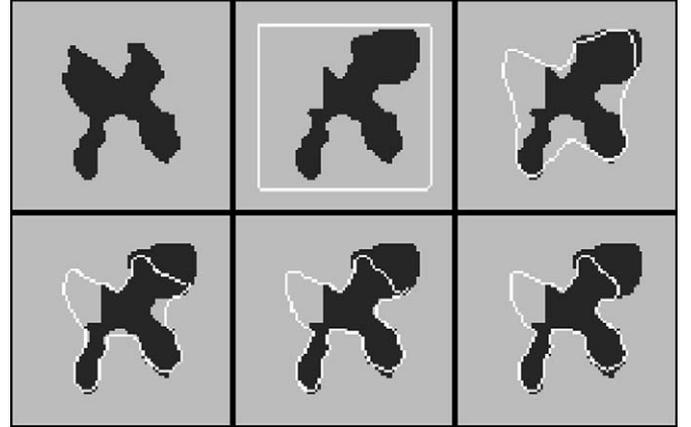


Fig. 7. From left to right, top to bottom: synthetic character. Some part was erased and an extra patch was added to the character. Four different stages of the curve evolution. This example also demonstrates the importance of the last term in Eq. (13). It attracts the contour to the region of permitted shapes.

composite shape prior. Their shape difference can be defined as

$$E_{shape} = \int (\Phi - \Phi_p)^2 dx dy \quad (9)$$

The corresponding gradient descent equation is given by

$$\frac{\partial \Phi}{\partial t} = 2(\Phi - \Phi_p) \quad (10)$$

The data term of our segmentation model is based on the Chan-Vese model with a regularization of the *geodesic* active contour, i.e.

$$E_{data} = E_{CV} + \alpha E_{GAC} \quad (11)$$

where E_{GAC} and E_{CV} are defined as in Eq. (4) and Eq. (6), respectively and α is a small constant such that the geodesic term on the right side is used only for regularization.

The incorporation of the shape term with the data term is based on the following observations:

- The prior confidence map η (see Section 2.2) determines whether we should follow the shape prior, i.e., the shape prior should be used only in regions where the η value is large.
- Regions with low η value indicate that none of the shape models was fitted to this region. In such cases there are two options, either that the real boundary passes in a different location, and information from the image is needed to attract the active contour to the correct place, thus the geodesic active contour should be dominant, or that the area is damaged and information from the image is not useful. In the latter case, we need a mechanism which will constrain the contour's evolution in order to remain as close as possible to a typical shape.

For this purpose we define a global shape constraint $\gamma(x, y)$ which acts as an adaptive force, and will constrain the evolution of the active contour by locally controlling the amount of allowed deformation. After aligning the shape models (represented by their SDF), we define $\mu(x, y)$ and $\sigma(x, y)$ to be the average and standard deviation of the aligned SDF, respectively. Regions with high σ values are allowed for more deformations, while regions with small σ values are more likely to remain close to the average μ . The function which limits the evolution of the active contour is defined as

$$\gamma(x, y) = \begin{cases} c & \text{if } \mu(x, y) + 2 \times \sigma(x, y) > 0 \\ -c & \text{if } \mu(x, y) - 2 \times \sigma(x, y) < 0 \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

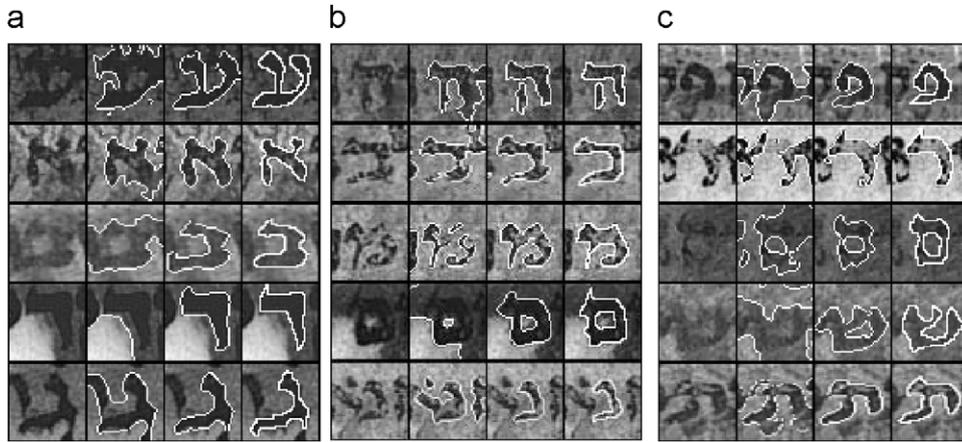


Fig. 8. (a)–(c) Segmentation results of different letter types. From left to right: original gray-scale character, segmentation without shape prior, segmentation with the proposed shape prior, and manual segmentation.

The meaning of this function is as follows. γ creates a narrow band of permitted deformations by adding a force outside the band (the constant c) which contracts the active contour, and a negative force which limits the contraction of the active contour by inflating it back. The final functional of the level set function Φ is composed both of the data term and the shape term:

$$E_{seg} = (1 - \eta)E_{data} + \eta E_{shape} + \gamma \quad (13)$$

and its corresponding level set equations are given by Eqs. (8) and (7) and $\gamma|\nabla\Phi|$, respectively. To demonstrate our approach, we corrupted the image shown in Fig. 7 by deleting part of it and adding an extra patch to its upper-right region. As can be seen, our model successfully segments the character by being attracted to the preserved boundaries while keeping the final contour close to the typical shape of the character.

5. Experimental results

We conducted our experiments on a set of four degraded documents from an antique Jewish prayer book composed by the famous scholar *Rav Seadia Gaon*, dated between the 11th and the 13th centuries. We tested our segmentation model on 500 degraded characters of 16 different Hebrew letters. For each type of letter, we extracted a set of 15 characters for model set creation (Section 2.3). All of the extracted characters were segmented manually in order to evaluate the segmentation results.

Our experiments were performed in two stages. First, we applied the recognition procedure as described in Section 3. As noted, integrating our shape prior into the NCC matching process, yielded 98% correct recognition. A careful examination of the misclassified characters showed that in most cases, without contextual knowledge of the word in which the character was embedded, recognition by human subjects involved high levels of uncertainty. An example of such character is depicted in Fig. 9.

In the second stage of our experiment, we tested the accuracy of our segmentation model both on characters that were correctly and incorrectly recognized. In all of our experiments, we used only four models for the shape prior construction. In order to evaluate our segmentation results, each one of the degraded characters was also segmented manually. Fig. 8 shows several examples of degraded character images along with the segmentation results with the proposed adaptive shape prior. Segmentation with the data term without prior (Eq. (11)) are shown in the second column, and the manual segmentation is shown in the last one. It is obvious that without

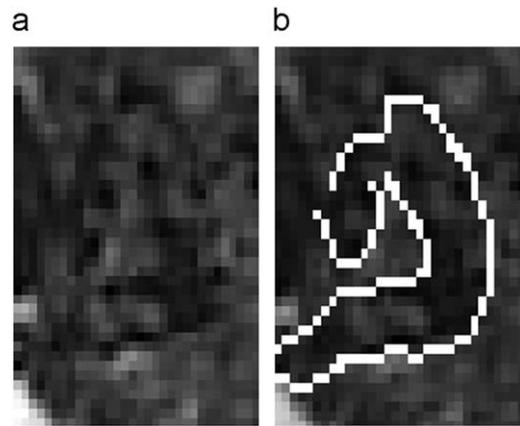


Fig. 9. (a) An example of misclassified character in the recognition process. (b) Despite of the misclassification, we tested whether our segmentation process would succeed if the recognition would have been correct. As can be seen, the results are very accurate.

Table 1
Evaluation of our segmentation model based on mean and maximum distance (in pixels) statistics.

Mean dist.	Std. dist.	Max. dist.	Std. max. dist.
0.73	0.8	3.8	2.25

The average size of the characters is 70 * 70 pixels.

incorporating the shape prior, accurate segmentation is far from being satisfactory. However, the proposed shape prior has achieved highly accurate results. To quantify the segmentation results, for each of the degraded characters in our set we measured the mean and maximum distance between its boundary and the boundary of the corresponding manually segmented character. The results are summarized in Table 1.

6. Conclusions and future work

In this paper, we propose a novel approach for incorporating an adaptive shape prior into the recognition and segmentation of highly degraded characters. The shape prior is constructed in a data driven manner, where fragments of characters are fitted to the segmented character. The resulting shape priors are integrated in a level set function. The influence of the shape prior is proportional to its degree of

local fitting to the degraded character. The proposed model achieves very accurate results both in recognition and segmentation of highly degraded characters.

There are two main directions in which the current work could be extended. In the current work, we tested our model on manually extracted single characters. Extending our work to whole documents, i.e., full recognition and segmentation would be of great value. In addition, we believe that the proposed approach of adaptive shape prior can be integrated in other recognition and segmentation schemes.

Acknowledgments

This work was partially supported by the Lynn and William Frankel Center for Computer Sciences, the Israel Ministry of Science and Technology, and by the Paul Ivanier Center for Robotics and Production Management at Ben-Gurion University, Israel. We would like to thank Dr. Uri Ehrlich from the Goldstein-Goren department of Jewish thought, Ben-Gurion University of the Negev, for supplying the documents.

References

- [1] B. Gatos, I. Pratikakis, J.S. Perantonis, Adaptive degraded document image binarization, *Pattern Recognition* 39 (2006) 317–327.
- [2] I. Bar-Yosef, Input sensitive thresholding for ancient Hebrew manuscript, *Pattern Recognition Lett.* 26 (8) (2005) 1168–1173.
- [3] M. Droettboom, Correcting broken characters in the recognition of historical printed documents, in: *Proceedings of the 3rd ACM/IEEE-CS Joint Conference on Digital Libraries, 2003*, pp. 364–366.
- [4] J.D. Hobby, T.K. Ho, Enhancing degraded document images via bitmap clustering and averaging, in: *Fourth International Conference Document Analysis and Recognition, 1997*, pp. 394–400.
- [5] B. Allier, N. Bali, H. Emptoz, Automatic accurate broken character restoration for patrimonial documents, *Int. J. Doc. Anal. Recognition* 8 (2006) 246–261.
- [6] C. Xu, J.L. Prince, Snakes, shapes, and gradient vector flow, *IEEE Trans. Image Process.* 7 (1998) 359–369.
- [7] T. Wakahara, Y. Kimura, A. Tomono, Affine-invariant recognition of gray-scale characters using global affine transformation correlation, *IEEE Trans. Pattern Recognition Mach. Intell.* 23 (2001) 384–395.
- [8] M. Oliveira, B. Bowen, R. McKenna, Y. Chang, Fast digital image inpainting, in: *Proceedings of the International Conference on Visualization, Imaging and Image Processing, 2001*, pp. 261–266.
- [9] A. Kass, M. Witkin, D. Terzopoulos, Snakes: active contour models, *Int. J. Comput. Vision* 1 (1988) 321–332.
- [10] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, *Int. J. Comput. Vision* 22 (1997) 61–79.
- [11] T. Chan, L. Vese, Active contours without edges, *IEEE Trans. Image Process.* 2 (2001) 266–277.
- [12] N. Paragios, R. Deriche, Geodesic active regions and level set methods for motion estimation and tracking, *Comput. Vision Image Understanding* 97 (2005) 259–282.
- [13] D. Mumford, J. Shah, Optimal approximations by piecewise smooth functions and associated variational problems, *Commun. Pure Appl. Math.* 42 (1989) 577–684.
- [14] R. Kimmel, Fast edge integration, in: S. Osher, N. Paragios (Eds.), *Geometric Level Set Methods in Imaging Vision and Graphics, 2003*, pp. 59–77.
- [15] S. Osher, J.A. Sethian, Fronts propagating with curvature dependent speed: algorithms based on Hamilton–Jacobi formulation, *J. Comput. Phys.* 79 (1988) 12–49.
- [16] Y. Chen, H.D. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K.S. Gopinath, R.W. Briggs, E.A. Geiser, Using prior shapes in geometric active contours in a variational framework, *Int. J. Comput. Vision* 50 (2002) 315–328.
- [17] M.E. Leventon, W.E.L. Grimson, O. Faugeras, Statistical shape influence in geodesic active contours, in: *CVPR, 2000*, pp. 316–323.
- [18] D. Cremers, N. Sochen, C. Schnorr, A multiphase dynamic labeling model for variational recognition-driven image segmentation, *Int. J. Comput. Vision* 66 (2006) 67–81.
- [19] T. Chan, W. Zhu, Level set based shape prior segmentation, in: *CVPR'05, vol. 2, 2005*, pp. 1164–1170.
- [20] D. Cremers, M. Rousson, R. Deriche, A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape, *Int. J. Comput. Vision* 72 (2007) 195–215.
- [21] N. Paragios, M. Rousson, V. Ramesh, Matching distance functions: a shape-to-area variational approach for global to local registration, in: *ECCV '02, 2002*, pp. 775–789.

About the Author—ITAY BAR-YOSEF is a PhD student at the computer science department of the Ben-Gurion University of the Negev, Israel. He received his B.Sc. and M.Sc. degrees from the Computer Science Department of the Ben Gurion University of the Negev, Israel in 2001 and 2003, respectively. His main research interests are in the area of historical documents analysis and image segmentation, variational methods, active contours, and level sets.

About the Author—ALIK MOKEICHEV is a PhD student at the computer science department of the Ben-Gurion University of the Negev, Israel. He received his B.Sc. degree from the Computer Science Department of the Hebrew University in 2001 and his M.Sc. degree from the Computer Science Department of the Ben-Gurion University of the Negev in 2006. His main research interests are computer vision and low level biological vision, statistical data analysis.

About the Author—KLARA KEDEM is a Professor at the Department of Computer Science at the Ben-Gurion University of the Negev. Her main research interests are computational geometry and bioinformatics.

About the Author—ITSHAK DINSTEIN received his BSc. in electrical Engineering from The Technion, Israel Institute of Technology, in 1965, and his MSc and PhD. degrees in Electrical Engineering from The University of Kansas in 1969 and 1974 respectively. He was with Comsat Labs, Clarksburg, Maryland, USA, during 1974–1977 working on digital television. He has been with Ben-Gurion University, Beer Sheva, Israel since 1977, where he is now a professor of Electrical Engineering. He has been the chairman of the Electrical and Computer Engineering Department from September 2004 to August 2007. He was a visiting scientist at IBM Research Laboratory in San-Jose, California, during 1982–1984, working on industrial visual inspection, and a visiting associate professor at Polytechnic University, Brooklyn, New-York, during 1988–1990. His current research interests include document understanding and analysis, and machine visual inspection.

About the Author—URI EHRLICH is the academic director of the Goldstein-Goren Department of Jewish thought at Ben-Gurion University. He is also the head of Mifal Hatefila—a national research project on the early version of Jewish Prayer founded by the Israel Academy of Sciences and Humanities.