NEURAL TECHNIQUES FOR OUTLINING THE LUNGS FROM MR IMAGES OF THE THORAX

S.J. GILSON, I. MIDDLETON and R.I. DAMPER

Image, Speech and Intelligent Systems Research Group,
Department of Electronics and Computer Science,
Building 1, University of Southampton,
Southampton, SO17 1BJ, UK
(emails: `{sjg95r|im94r|rid}@ecs.soton.ac.uk`)

ABSTRACT

Two of the most common problems in computer vision are the detection of relevant edges and linking them to extract object shape. In this paper, we describe a two-stage neural system which segments lung outlines from transverse magnetic resonance images of the thorax. The first stage uses a multilayer perceptron (MLP) to derive an edge-point image. It is trained (on a small proportion of the total MRI data) to produce a lung-boundary or non-lung-boundary classification of each pixel. The MLP has no notion that there are two lungs each with a closed contour: every pixel is processed entirely independently. In the second stage, therefore, a connectionist-based deformable model is applied to link the edges produced in the first stage. We describe here an implementation of the system, with discussion of results and potential use.

KEYWORDS: neural networks, image segmentation, shape extraction.

INTRODUCTION

Automatic segmentation of images and feature extraction are important computer-vision goals. In medical applications, image quality is often compromised because of practical needs to limit acquisition time, cost, etc. Also, the extreme complexity of the human body means that isolating specific regions implies obvious problems, so that matching organs with simple geometric constructs is difficult or impossible.

We present an automatic technique for extracting the lung boundaries from magnetic resonance (MR) images of the thorax. This was motivated by a requirement from Southampton General Hospital, UK, to register the MR images with SPECT images of the thorax. The MR images comprise a set of approximately 30 transverse slices per subject. Figure 1(a) shows a typical slice (number 12, subject AC) in which the lungs are clearly visible as two large, low-intensity regions within the torso. They are similar in intensity to the background and blood vessels because both air and moving blood yield virtually no MR signal. Each slice is composed of (256×256) pixels corresponding to the $T1$-weighted (see [1]) values of each scan, and is accompanied by a corresponding ('target') image containing just the outline of the lungs (figure 1(b)). Target images were obtained using a semi-automatic technique based on region growing which is both time-consuming and tedious: it is the aim of this work to replace this with a fully-automatic procedure. Note that these outline images provide the ‘gold standard’ for both network training and evaluation.

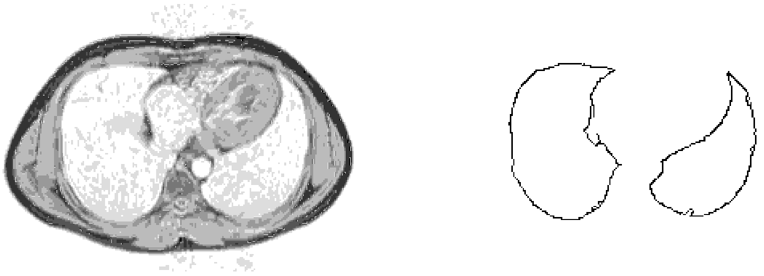

(a) (b)

Figure 1. (a) Typical slice of MRI data from a human male thorax (grey-scale has been inverted to show detail. (b) Target data produced by semi-automated method.

The segmentation procedure is a two-stage process (see figure 2). First, using a multi-layer perceptron (MLP), each pixel of each slice is classified into *lung-boundary* and *non-lung-boundary* pixels [2, 3]. This classification is not perfect: the two lung boundaries are not always closed and a number of non-lung-boundary pixels (mainly from the boundaries of the great blood vessels) are misclassified. Hence, the second stage employs a deformable model to produce a closed, noise-free contour for each lung from the above classification [4]. The deformable model also produces a coordinate-based representation of the lung boundaries which is more versatile for manipulation and visualisation than the bitmaps obtained from the semi-automatic method, or from the MLP.

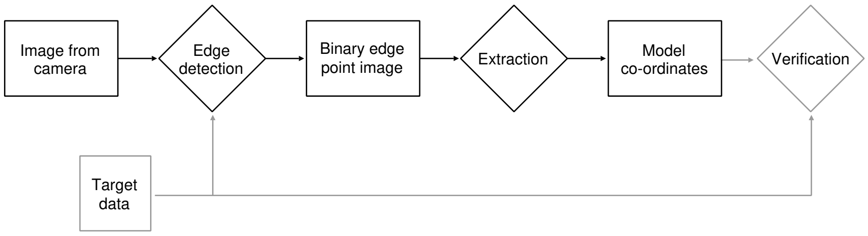

Figure 2. Flow of data required for extraction from binary edge point data. The target data are used in training the MLP (edge detection) and in verification.

In this paper, we describe enhancements to both the classification and extraction stages that result in substantial improvements in the robustness of the segmentation.

CLASSIFICATION

The inputs to the MLP are the intensity values of the pixels from a (7×7) window centred on the pixel to be classified. This window size was found to be the smallest that enabled the lung boundary to be distinguished from the great blood vessel boundaries. The output is a single node trained to have an activation of 1.0 for an input window centred on a lung boundary, and 0.0 otherwise. The network has a single hidden layer of 30 nodes.

The network was trained using error back-propagation [5, 6] with weight elimination [7] (to improve the network's generalisation ability). The training data were constructed iteratively: a proportion of the misclassified examples were added to the training set and used for retraining. The process is initiated from a small random selection of lung-boundary and non-lung-boundary examples, and is terminated when a reasonable classification (on a given slice) is achieved. In this way, the training data for a given problem can be constructed automatically provided good labeled data already exist.

Results using two training schemes are used in this paper. The first uses data from multiple slices of a single subject. This enables the network to classify accurately the slices on which it was trained: see figure 3(a). The network can also classify to a usable degree (see next section) the remaining slices from that subject. However, the performance on slices from other subjects is very poor (figure 3(b)) but can be improved by training the network on additional slices from other subjects. Thus, the second training scheme uses multiple slices from three subjects.

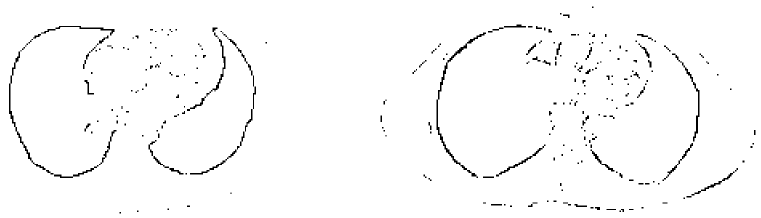

(a) (b)

Figure 3. MLP output after training on every third slice from AC. (a) Slice 12 from subject AC. (b) Slice 12 from subject NH. Note that, despite the higher levels of noise in (b), the lung boundaries are still the dominant feature in the image. Each image contains approximately 600 pixels.

The MLP classifies each pixel independently of the others, and therefore has no notion of a circular closed contour. Consequently, the boundaries it produces are often fragmented and noisy (false negatives and false positives respectively). We now use a deformable model to link the boundary segments together, while attempting to ignore noise.

EXTRACTION

We use a development of a deformable model technique to link the lung boundary segments together. The model consists of many simple units, (c.f. neurons) each connected to its neighbours. By applying simple attraction (c.f. weights) and regularising forces, the units converge toward dense regions of the image – in this case the lung boundaries.

The elastic net [8, 9] was originally designed for finding near-optimal solutions to the travelling salesman problem, but has also found application in edge linking binary images [10, 11]. However, the travelling salesman problem also implies a noise-free environment, and, thus, the elastic net has no tolerance for noise. Despite this, the technique does have some desirable properties as a generic shape extraction method; most notably that it considers all the information in the image as it converges (i.e. a *global* technique). Contrast this with active contours ('snakes') [12], where each unit of the snake only considers a small local neighbourhood, and usually requires expert seeding and guidance. In practice, this means that the elastic net, if adapted to tolerate noise, would make a good choice for arbitrary shape extraction, without requiring expert initialisation.

We discuss here further enhancements to a noise-tolerant elastic net algorithm [4] which enables improved performance through consideration of three-dimensional information. Since a spatial sequence of MRI scans is available, extra noise tolerance can be found by extracting the lung boundaries in parallel, using one net per slice.

Each of the s nets is connected to its two immediate neighbours by an inter-layer tension force, whose strength is governed by a parameter γ. Now, each of the units, $\mathbf{u}_j$, is subjected to: attraction forces from edge points, $\mathbf{p}$, intra-layer forces from its two neighbours, and the new inter-layer forces from its correspondents in the neighbouring nets. The update rule is now a two-step process (which reduces instabilty because of the direct competition between the three forces):

$$\begin{aligned}\left(\Delta\mathbf{u}_{j,l}^{t+1}\right)' &= \alpha\sum_{i=1}^{N} G_{ij}\left(\mathbf{p}_{i,l}-\mathbf{u}_{j,l}^{t}\right)+K\beta\left(\mathbf{u}_{j+1,l}^{t}-2\mathbf{u}_{j,l}^{t}+\mathbf{u}_{j-1,l}^{t}\right)\\ \Delta\mathbf{u}_{j,l}^{t+1} &= K\gamma\left\{\left(\mathbf{u}_{j,l+1}^{t+1}\right)'-2\left(\mathbf{u}_{j,l}^{t+1}\right)'+\left(\mathbf{u}_{j,l-1}^{t+1}\right)'\right\}\end{aligned}$$

where a new index, l, denotes the current layer of the multi-layered data. G_{ij} is the normalised αcut Gaussian attraction function, and α is the 'learning rate' of the net. β governs intra-layer tension strength, and K is a simulated annealing term, controlling convergence.

Each net is initialised as a large lung-boundary encompassing circle, which, when viewed in three dimensions, appears as a cylinder enveloping the boundaries. As in the 2D case, the cylinder is attracted toward image features, gradually deforming until each net converges on the boundary in its slice. The nets now represent the surface of the object (see figure 4).

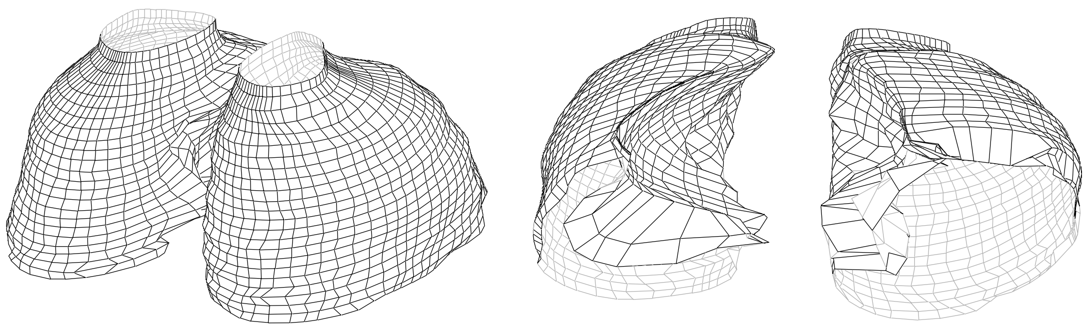

Figure 4. Two wire-frame projections of the same lungs (subject AC, 24 slices). Note the large lung deformations in the right hand image, caused by the heart and liver.

We quantify the performance of the modified elastic net using the *effectiveness* measure [13]. Although the goal is to extract the boundary, we evaluate this using the *regions* defined by the extracted and the target contours – since we are trying to model the target with an approximate contour, this evaluation scheme reduces the significance of small pixel-sized deviations between extracted and target contours. Thus, we consider the pixels contained in the boundary of the extracted and target outlines, and by examining their intersection, we obtain a similarity measure, E:

$$precision = \frac{N_{intersection}}{N_{result}} \qquad recall = \frac{N_{intersection}}{N_{target}}$$

$$E = 1 - \frac{precision \times recall}{\alpha \times recall + (1-\alpha) precision} \qquad 0 \le E < 1,\ 0 \le \alpha \le 1$$

Here, α can be used to bias precision or recall, which in this application, correspond to a model which is bigger or smaller than the target outline. Since both of these situations are equally undesirable, we use $\alpha = 0.5$. E now represents an inverse measure of goodness.

The MLP was trained using selected slices from subjects AC, CB and LP. All slices from these three subjects, and an 'unseen' subject, NH, were then classified, and the resulting edge segments linked by the modified elastic net. The E-values are then evaluated for each slice of each subject, and an average score is shown in table I – this reflects the overall performance of both the classification and extraction stages. To put these scores into context, extracting the lung boundaries from the semi-automatically extracted outlines (noise-free) produced an E-value of 0.06. The performance of the modified elastic net on the MLP's output is reduced because of its inability to cope with the dramatic deformation around the liver in the lower slices of the data set. This is usually a result of poorly classified boundary segments on these slices.

TABLE I. Segmentation performance for four subjects.

	N_{target}	N_{result}	$N_{intersection}$	precision	recall	E
AC	5504	5789	5111	0.87	0.93	0.11
CB	7127	8263	6722	0.83	0.93	0.14
LP	9613	10376	9120	0.85	0.93	0.12
NH	6709	7696	6500	0.84	0.97	0.10

Two problems arise from integrating information from neighbouring slices in the manner used by the modified elastic net. First, the elasticity of the net favours smooth boundary changes between successive slices. This problem manifests itself in the lower slices of the thorax, where the liver forces a dramatic deformation of the lungs, and, because of the coarse resolution of the MRI images (approximately 8mm between slices), the resulting images show large changes in boundary position, rather than smooth transitions. This largely accounts for the poor segmentation performance, especially for subjects AC and CB.

Secondly, although increasing tolerance to uncorrelated noise, the integration of 3D information reduces tolerance to correlated noise, as might be found from consistent incorrect classification of major blood vessels across successive slices. Under these conditions, the extraction stage may incorporate the blood vessel into the extracted outline, producing an incorrect segmentation.

CONCLUSIONS

We have described extensions to two stages of an automated lung-boundary extraction system for processing magnetic resonance images of the human torso. The grey-scale MRI scans are first classified into boundary and non-boundary pixels using an MLP. Since the classification is imperfect, we then use a deformable model (one per lung) to extract the complete boundary, producing a contour description of the lung. These can then be combined to produce a full 3D model of the lungs, which could be of use for surgery planning and/or diagnosis.

The extensions have shown an improved level of robustness when compared to earlier levels of development ([4, 2]). While performance is still imperfect, we anticipate that future work will improve both the quality of the classification, and the ability of the modified elastic net to cope with the deformations around the liver, resulting in a reliable, automated extraction method.

Acknowledgement

Author SJG is supported by a Research Studentship from the UK Engineering and Physical Sciences Research Council (EPSRC). Author IM is supported by a Cooperative Award in Science and Engineering from the EPSRC and Roke Manor Research, Romsey, UK.

REFERENCES

1. Webb, S. *The Physics of Medical Imaging*. Hilger, Bristol, 1988.
2. Middleton, I., and Damper, R. "Segmentation of magnetic resonance images of the thorax by back-propagation". In *IEEE International Conference on Neural Networks* (Perth, Western Australia, 1995), vol. 5, pp. 2490–2494.
3. Middleton, I., Damper, R. I., Moore, E. A., and Fleming, J. S. "Lung boundary identification in magnetic resonance images using feedforward neural networks". *IEEE Transactions on Medical Imaging* (1996). Submitted for publication.
4. Gilson, S. J., Middleton, I., and Damper, R. I. "A localised elastic net technique for lung boundary extraction from magnetic resonance images". In *IEE 5th International Conference on Artificial Neural Networks* (July 1997), no. 440, pp. 199–204.
5. Rumelhart, D. E., Hinton, G., and Williams, R. "Learning representations by back-propagating errors". *Nature* 323 (1986), 533–536.
6. Chauvin, Y., and Rumelhart, D., Eds. *Backpropagation: Theories, Architectures and Applications*. Lawrence Erlbaum, Hillsdale, NJ, 1995.
7. Weigend, A., Rumelhart, D., and Huberman, B. "Generalisation by weight-elimination with application to forecasting". *Advances in Neural Information Processing* 3 (1991), 875–882.
8. Durbin, R., and Willshaw, D. "An analogue approach to the travelling salesman problem using an elastic net method". *Nature* 326 (April 1987), 689–691.
9. Durbin, R., Szeliski, R., and Yuille, A. "An analysis of the elastic net approach to the travelling salesman problem". *Neural Computation* 1 (1989), 348–358.
10. Abrantes, A. J., and Marques, J. S. "Unified approach to snakes, elastic nets, and Kohonen maps". *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 95* 5 (May 1995), 3427–3430.
11. Gilson, S., and Damper, R. "An emprical comparison of neural techniques for edge linking of images". In press (*Neural Computing and Applications*).
12. Kass, M., Witkin, A., and Terzopoulos, D. "Snakes: Active contour models". In *Proceedings of 1st International Conference on Computer Vision* (London, June 1987), pp. 259–268.
13. van Rijsbergen, C. *Information Retrieval*, second ed. London: Butterworth, 1979.