
G5.5 Hopfield neural nets for the optimum segmentation of medical images

R Poli

School of Computer Science
The University of Birmingham (UK)

G Valli

Department of Electronic Engineering
University of Florence (Italy)

Abstract

In this chapter we present a general purpose neural architecture for segmenting 2-D and 3-D medical images. The architecture is based on a continuous *Hopfield neural network* including one or more sets of 2-D layers of neurons with local connections. This architecture can be specialised to perform the *segmentation of 2-D images*, the multi-scale segmentation of 2-D images and the segmentation of 3-D images by simply changing the number of such sets and/or the size of the component layers. By changing synaptic weights the architecture can adapt to the differences existing between tomographic and radiographic images. The segmentation produced by this architecture is optimum with respect to a goodness criterion which establishes the tradeoff between sensitivity and robustness. The chapter describes the derivation of the architecture and some experimental results obtained with synthetic and real medical images.

C1.3.1

G1

Keywords: Hopfield Networks, Image Analysis, Image Segmentation, Medical Images.

G5.5.1 Introduction

The general objective of the *segmentation* of medical images is to find regions which represent single anatomical structures. The availability of such regions not only makes tasks such as interactive visualisation and automatic measurement of clinical parameters directly feasible but is also the starting point to use more sophisticated computer vision techniques and perform higher-level tasks such as 3-D shape comparison and recognition (Poli *et al.* 1994).

G1

Unfortunately, due to the presence of image noise, masking structures, biological shape variability, tissue inhomogeneity, imaging-chain anisotropy and variability, etc. the segmentation of medical images is a very hard problem. Therefore, to obtain reliable segmentation algorithms almost invariably researchers have been obliged to exploit as much *a priori* information as possible.

Knowledge of statistical properties of the gray levels of the image is a kind *a priori* information that has been extensively exploited in the case of Magnetic Resonance (MR) and Computed Tomography (CT) images (see for example Raya 1990, Lei and Sewchand 1992, Gerig *et al.* 1992, Amatur *et al.* 1992, Ozkan *et al.* 1993). Despite the differences existing among these methods, they share the idea of considering each pixel as a separate entity to be classified, thus neglecting the spatial correlation between measurements due to cohesion of matter.

Spatial correlation is considered as more important in other methods, such as those based on mathematical morphology operators (Higgins *et al.* 1990, Klingler *et al.* 1988, Thomas *et al.*

1991, Joliot and Mazoyer 1993), on rule-based expert systems (Catros and Mischeler 1988, Manos *et al.* 1993, Li *et al.* 1993), on special purpose computer vision techniques (Raman *et al.* 1993, Coppini *et al.* 1993, Deklerck *et al.* 1993) or on neural nets trained with the *backpropagation algorithm* (Silverman and Noetzel 1990, Toulson and Boyce 1992, Coppini *et al.* 1993). However, in addition to the spatial correlation between measurements all these methods exploit another kind of *a priori* information: the anatomical knowledge about which structures are present in the image, where they usually are, what they usually look like, etc. If on the one hand this information improves considerably the robustness of segmentation algorithms, on the other hand it drastically reduces their generality and their applicability to different kinds of images or anatomical districts. Therefore, anatomical-information-based methods do not seem good candidates to build general purpose segmentation systems for medical images.

C1.2

To overcome these problems and build a general-purpose segmentation system for medical images, we adopted a different approach inspired by biological vision.

G5.5.2 Approach and Objectives

Vision is ruled by principles, such as perceptual grouping, selection and discrimination, which mostly depend on regularities of nature such as cohesiveness of matter or existence of bounding surfaces (Marr 1982, Reuman and Hoffman 1986). As these properties are valid also for the anatomical structures present in medical images, they can be exploited to build segmentation systems for such images. If no other source of information is used, the resulting segmentation algorithms are independent of the imaging modality, of the scanning parameters, of the imaged district, etc. and, therefore, can be used for general-purpose medical-image segmentation.

Regularities of nature can be exploited in a very simple way by using grouping or discrimination criteria based for example on the idea that pixels which are close to each other and have similar gray levels have a high probability of representing the same object and, therefore, should be grouped together. However, even if the strategy is simple, in order to design a general purpose segmentation algorithm for medical images a number of requirements must be met which can make the actual implementation of the strategy quite complex. Let us analyse these requirements:

- The segmentation algorithm should be maximally sensitive to small structures or to structures with a low contrast (possible lesions or tumours in early stages).
- The algorithm should be maximally robust with respect to the noise, texture and slow intensity changes typically present in medical images.
- The algorithm should be able to adapt to the differences existing among the processes of generation of images obtained from different imaging devices. Therefore, it should be able to process not only 2-D tomographic images but also 3-D and X-ray projective ones.
- A segmentation algorithm to be integrated in more complex analysis systems should be able to perform multi-scale segmentation as, in many applications, segmentation is analysed by multiple modules requiring different levels of detail.
- An algorithm to be used with imaging devices (e.g. cine-CT scanners) which can produce hundreds of images per patient should be suitable for parallel, high-speed implementation.

The first two requirements counteract each other and, therefore, any segmentation algorithm can only produce results that represent a trade-off between them. In order to achieve optimum compromises it is first necessary to define a quantitative criterion of goodness of segmentation which takes sensitivity and robustness into account, and then to optimise it for any specific image. Therefore, the problem of medical image segmentation can be seen as a problem of combinatorial optimisation.

Unfortunately, for any given image the space of possible solutions to this optimisation problem is huge and conventional optimisation techniques tend to fail on it. Therefore, following recent approaches in the field of natural scene segmentation (Darrell *et al.* 1990, Reed 1992, Wang *et al.* 1992), we decided to solve it by using an architecture based on continuous Hopfield neural nets (Hopfield 1984), a computational paradigm which can effectively search huge solution spaces.

Hopfield nets can be seen as dynamical systems which tend to relax into states of which minimise

C1.3.1

the following energy function:

$$E_{net} \approx -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N T_{ij} v_i v_j - \sum_{i=1}^N i_i v_i. \quad (\text{G5.5.1})$$

where v_i is the output of neuron i , i_i is its external input and T_{ij} is the weight of the connection from neuron j to neuron i . Thanks to this minimum-seeking behaviour, Hopfield nets can be used to solve optimisation problems (Hopfield and Tank 1985, Hopfield and Tank 1986). The basic strategy is as follows: a) to preprocess, when needed, the input data, b) to find a binary representation for the solutions of the problem so that they can be mapped into the stable states of the neurons of a Hopfield net, c) to define a quadratic (symmetric) energy function whose minimisation leads to an optimum solution of the problem and then calculate weights and external inputs, d) to initialise and let the network relax into a stable state to be then mapped back into a solution for the original problem.

In the following we describe how these steps, applied to the problem of medical-image segmentation, lead to an architecture that not only provides the optimum sensitivity/robustness trade-off but also meets the other requirements listed above.

G5.5.3 Segmentation of Tomographic Images

In this case the input data of the segmentation algorithm is a 2-D tomographic image denoted with the symbol $I(x, y)$. Normally this data needs no preprocessing and, therefore, the first step for solving the segmentation problem is finding a binary representation for its solutions.

Binary Representation

We adopted a representation suggested by the analogy of the segmentation process with that of colouring geographic maps Bilbro *et al.* 1987. This analogy indicates that, in order to represent the regions (“states”) obtained from the segmentation of an image, only a reduced number of labels (“colours”) are needed, as long as different labels are associated to connected regions (“bordering states”). Therefore, as shown in Figure G5.5.1, a segmentation can be represented with a small set of 2-D layers of neurons, each layer represent different a label.




Figure G5.5.1. Synthetic 8×8 image (left), a possible labelling with 4 colours (centre), and the related binary representation with 4 layers of neurons (right). (Active neurons are represented as filled circles.)

Energy Function

The next step is the definition of a quadratic energy function E_{net} whose minimisation gives an optimal solution to the segmentation problem. We adopted an energy function partially inspired by the one suggested in (Hopfield and Tank 1985, Hopfield and Tank 1986) for the solution of the travelling-salesman problem and the one proposed in (Bilbro *et al.* 1987) for the segmentation of

signals with simulated annealing. As the fixed points of Hopfield nets tend to be the vertices of the hypercube $[0, 1]^N$, we were able to design E_{net} on the hypothesis of binary neurons, i.e. $v_{xyc} \in \{0, 1\}$.

E_{net} includes two parts: the *syntax energy* E_{syntax} which enforces the syntactic correctness of the solutions (i.e. prevents the network from settling into non-binary states or states which cannot be mapped back to solutions of the segmentation problem), and the *semantics energy* $E_{semantics}$ which is our criterion of goodness of segmentation. The two parts are added so that $E_{net} = E_{syntax} + E_{goodness}$.

Syntax Energy The syntactic correctness of the solutions requires that *one and only one neuron is active among the neurons which represent a given pixel*, i.e. $\exists!c : v_{xyc} = 1$. This constraint can be enforced by including in E_{syntax} terms such as $\sum_{c_1} \sum_{c_2 \neq c_1} v_{xyc_1} v_{xyc_2}$ and $(\sum_c v_{xyc} - 1)^2$ (the latter prevent the network from settling in the non-valid null solution $v_{xyc} = 0, c = 1, 2, \dots$). By summing up these terms for all the pixels in the image we obtain:

$$E_{syntax} = \frac{K_1}{2} \sum_x \sum_y \sum_c \sum_{\hat{c} \neq c} v_{xyc} v_{xy\hat{c}} + \frac{K_2}{2} \sum_x \sum_y \left(\sum_c v_{xyc} - 1 \right)^2 \quad (\text{G5.5.2})$$

where K_1 and K_2 are constant values.

Semantics Energy The goal of the semantics energy is that of driving the network towards segmentations that represent an optimum compromise between sensitivity and robustness. Therefore the semantic energy includes two terms, the *sensitivity energy* $E_{sensitivity}$ and the *robustness energy* $E_{robustness}$, which are summed up to give $E_{semantics} = E_{sensitivity} + E_{robustness}$.

Sensitivity Energy The sensitivity energy should force the network to perform a segmentation revealing any transition between different tissues, i.e. any change in the image gray-levels. In order to obtain this effect, $E_{sensitivity}$ must include terms which increase when neighbouring pixels lying across a boundary have the same label. We used terms such as $\sum_c v_{xyc} v_{\hat{x}\hat{y}c} \frac{dI(x, y)}{d\vec{n}(x, y, \hat{x}, \hat{y})}$ where $\vec{n}(x, y, \hat{x}, \hat{y}) = \frac{(\hat{x}, \hat{y}) - (x, y)}{\|(\hat{x}, \hat{y}) - (x, y)\|}$, and (x, y) and (\hat{x}, \hat{y}) are neighbouring pixels. These terms must be present for all pixels lying in a neighbourhood \mathcal{B}^{xy} which does not contain pixels too close to or too far from (x, y) . (We adopted the simplest neighbourhood satisfying these requirements: $\mathcal{B}^{xy} = \left\{ (\hat{x}, \hat{y}) \mid 2 \leq \sqrt{(\hat{x} - x)^2 + (\hat{y} - y)^2} \leq 2\sqrt{2} \right\}$). Thus, the complete expression of the sensitivity energy is:

$$E_{sensitivity} = \frac{K_4}{2} \sum_x \sum_y \sum_{(\hat{x}, \hat{y}) \in \mathcal{B}^{xy}} \sum_c v_{xyc} v_{\hat{x}\hat{y}c} \frac{dI(x, y)}{d\vec{n}(x, y, \hat{x}, \hat{y})} \quad (\text{G5.5.3})$$

where K_4 is a constant value.

Robustness Energy The aim of $E_{robustness}$ is to reduce the effects of noise and texture. Since noise and texture tend to produce very small regions, $E_{robustness}$ should favour the construction of large regions which have a high probability of representing single anatomical structures. This can be obtained using the constraint: *pixels which are close to each other should have the same label*. The constraint can be implemented using terms of the form $-\sum_c v_{xyc} v_{\hat{x}\hat{y}c}$, for all the pixels (\hat{x}, \hat{y}) in a 4-connected neighbourhood \mathcal{N}^{xy} of any given pixel (x, y) . The total robustness energy results to be

$$E_{robustness} = -\frac{K_5}{2} \sum_x \sum_y \sum_{(\hat{x}, \hat{y}) \in \mathcal{N}^{xy}} \sum_c v_{xyc} v_{\hat{x}\hat{y}c} \quad (\text{G5.5.4})$$

where K_5 is a constant value.

Once E_{net} is defined, the weights and the external inputs of the network can be computed easily (for example by comparing the expression of E_{net} with the left-hand side of Equation G5.5.1).

Network Initialisation

Hopfield network can be simulated by simply integrating numerically their motion equation until a stable state is reached. However, before doing that the state of the net has to be initialised. As the standard random initialisation method gives, in the present case, poor segmentation results, we adopted the strategy suggested in (Chen *et al.* 1991) which consists of initialising the network in an area of state space where a good solution is present. In this way the network has only to improve the solution instead of looking for it in the whole state space. As initial solution we used the segmentation produced by the following algorithm:

1. Let $I_{max} = \max_{x,y} I(x, y)$ and $I_{min} = \min_{x,y} I(x, y)$.
2. For each pixel (x,y) do:
 - 2.1. Let \hat{c} be the nearest integer which is less or equal to $\frac{I(x,y)-I_{min}}{I_{max}-I_{min}} + 1$.
 - 2.2. For each colour $c = 1, \dots$ do:
 - 2.2.1. Let $v_{xyc} = \begin{cases} \alpha & \text{if } c = \hat{c} - 1 \text{ or } c = \hat{c} + 1 \\ 1 - 2\alpha & \text{if } c = \hat{c} \\ 0 & \text{otherwise} \end{cases}$

Extensions to 3-D and multiscale segmentation

The extension of the method to the segmentation of 3-D images can easily be obtained by introducing 3-D neighbourhoods and 3-D image derivatives as well as by adding an extra sum in Equations G5.5.2, G5.5.3 and G5.5.4.

The extension to multi-scale segmentation requires a preprocessing step as segmentation has to be performed simultaneously on multiple, smoothed and decimated versions of the original image. Such images, denoted with the symbol $I(x, y, s)$, are build recursively from one another according to the formula:

$$\begin{cases} I(x, y, 1) = I(x, y) \\ I(x, y, s + 1) = \frac{1}{4} \sum_{i=0}^1 \sum_{j=0}^1 I(2x + i, 2y + j, s) \quad \text{for } s=1, 2, \dots \end{cases}$$

After preprocessing, the various components of the energy function can be separately defined for each scale and summed up. However, in order for the segmentation performed at a given scale to influence and to be influenced by the segmentation being performed at other scales additional energetic terms such as $-v_{xycs}v_{(x/2)(y/2)c(s+1)}$ and $-v_{xycs}v_{(2x+i)(2y+j)c(s-1)}$ (for $i = 0, 1$ $j = 0, 1$) are needed.

Derivation of weights and inputs, initialisation and relaxation are performed as in the case of 2-D segmentation.

G5.5.4 Segmentation of X-ray images

The general criteria of goodness of segmentation introduced in the previous sections are valid also for projective X-ray images. However, the peculiarities of the physical process of generation of this kind of images impose a few changes.

Preprocessing

The approximate linearisation of the image generation process is a preprocessing step needed for X-ray image segmentation. This is obtained by performing an appropriate logarithmic transformation of the gray levels of the original image after which we can express

$$I(x, y) = \int_0^{d(x,y)} \mu(x, y, z) dz,$$

where $\mu(x, y, z)$ is the linear absorption coefficient of the tissue at coordinates (x, y, z) and $d(x, y)$ the thickness of the body in (x, y) . As any anatomical district contains a discrete number of structures

of interest, if we denote with $d_i(x, y)$ the thickness of the i -th structure and with μ_i the absorption coefficient of such a structure, we can rewrite

$$I(x, y) = \sum_{i=1}^N \mu_i d_i(x, y). \quad (\text{G5.5.5})$$

Binary Representation

Anatomical structures which are overlaid or inside one another are represented by the same pixels in an X-ray image and, therefore, regions are no more constrained to form a tessellation of the image but can overlap.

To represent in binary form a segmentation with overlapping regions we adopted a set of 2-D layers of neurons like those used for the segmentation of tomographic images, with the important difference that each layer does not represent a different “colour” but a different anatomical structure.

Energy Function

Syntax Energy The syntactic correctness of solutions does not require any more that one and only one neuron inside the set of neurons which represent a given pixel be active, as this would mean that a pixel cannot represent more than one anatomical structure. However, syntax requires that, in stable states, each neuron of the network be completely excited ($v_{xyc} = 1$) or inhibited ($v_{xyc} = 0$). To obtain this effect we used a term of the form $v_{xyc}(1 - v_{xyc})$ for each neuron. As a result:

$$E_{syntax} = \frac{K_1}{2} \sum_x \sum_y \sum_c v_{xyc}(1 - v_{xyc})$$

where K_1 is a constant value.

Sensitivity Energy The function of $E_{sensitivity}$ is maximising the consistency of segmentation with respect to the image gray levels expressed by Equation G5.5.5. Unfortunately, to obtain a quadratic $E_{sensitivity}$ we had to add the hypothesis (only approximately valid) that the thickness of the structures shown in the X-ray image is constant, i.e. $d_i(x, y) = d_i$. On this hypothesis we can define the quantity $D_i = \mu_i d_i$ (estimated on the basis of the typical density and thickness of the structures of interest) and express $I(x, y) = \sum_c v_{xyc} D_c$. To force the network to settle into solutions (approximately) consistent with this equation we defined

$$E_{sensitivity} = \frac{K_2}{2} \sum_x \sum_y \left(\sum_c v_{xyc} D_c - I(x, y) \right)^2,$$

K_2 being a proper constant value.

Robustness Energy The robustness energy for X-ray image segmentation includes the same terms as in Equation G5.5.4. Unfortunately, in this case these terms alone can induce the diffusion of the activation of the neurons representing a given structure outside the boundaries of that structure. This happens because $E_{sensitivity}$ does not include any terms which force the neurons of a region to change their state in proximity of the boundaries of the structure represented by that region. This can be overcome by including also the constraint: *if a structure is not present in a given pixel, it is not present also nearby*. The resulting robustness energy turns out to be

$$E_{robustness} = -\frac{K_3}{2} \sum_x \sum_y \sum_{(\hat{x}, \hat{y}) \in \mathcal{N}^{xy}} \sum_c v_{xyc} v_{\hat{x}\hat{y}c} - \frac{K_4}{2} \sum_x \sum_y \sum_{(\hat{x}, \hat{y}) \in \mathcal{N}^{xy}} \sum_c (1 - v_{xyc})(1 - v_{\hat{x}\hat{y}c}).$$

where K_3 and K_4 are constant values.

Weights and external inputs can be easily obtained in the standard way. In order to ensure the convergence of the network to good solutions, also in this case we initialised it to a point of state space which represents a good segmentation. The initialisation algorithm is similar to that used for tomographic images.

G5.5.5 Experimental Results

The networks described in the previous sections have been tested both on synthetic images and on real tomographic and X-ray ones.

Synthetic images were generated by simulating the operation of a real tomographic device on an ellipsoidal organ surrounded by a homogeneous tissue. In order to test the robustness of the method, in addition to the blurring caused by the finite thickness of the slices (partial-volume effect), Gaussian white noise with zero mean and increasing standard deviation σ was included in the images. The resulting images were segmented using both the single-scale and multi-scale networks described in the previous sections and then compared with the exact segmentation obtained manually with images in which noise and partial-volume effect were absent. Table G5.5.1 shows the average errors obtained in these experiments for several different values of σ and for 1–4 interacting scales.

Noise σ	Scales			
	1	2	3	4
0	1.05	1.20	1.49	1.61
5	1.17	1.68	1.81	1.81
10	1.32	1.93	2.05	1.95
20	1.46	1.95	1.98	1.93
40	16.38	4.59	4.17	4.13
80	52.88	51.81	53.71	55.59

Table G5.5.1. Segmentation of synthetic tomograms: wrong assignments (%) vs. noise standard deviation and number of interacting scales.

The table reveals that, in the presence of noise with relatively small standard-deviation, there are no advantages in using multi-scale segmentation. Actually, for $\sigma = 0 - 20$, using 2–4 scales produces 14% to 55% more wrong assignments than in the single-scale case. However, in the presence of noise of higher intensity ($\sigma = 40$) multi-scale segmentation is much more reliable than single-scale one. Results are not satisfactory only when noise standard-deviation is extremely high ($\sigma = 80$).

The accuracy shown by the method in the experiments with synthetic images has been confirmed by numerous experiments with real tomograms. For example, Figure G5.5.2 illustrates how, in segmenting an MR image of the thorax, the net has correctly identified most of the anatomical districts of clinical interest (e.g. lungs, sub-cutaneous fat, muscular tissue, right atrium, right ventricle, backbone and pulmonary artery). Another example is represented by Figure G5.5.3 which shows an MR slice of the head along with the multiscale segmentation produced by the net. Segmentation has been performed jointly at three different scales: 128×128 , 64×64 and 32×32 . At the lowest resolution there are only eight regions the largest five of which represent the most significant anatomical structures: white matter, gray matter, cerebro spinal fluid (CSF) in the ventricles, fat with bone, and background. These regions can be easily recognised and used to guide a complete interpretation of the image. At 64×64 resolution the boundaries of white matter, gray matter and CSF become more complex and new regions are present to represent the difference between fat and bone and between thin and thick areas of the ventricles. Maximum accuracy is reached at the highest resolution where, despite noise and texture the most important structures are still represented by a single or a small number of large regions.

The method has also been tested on X-ray images. For example, Figure G5.5.4 (left) shows a cine-angiographic X-ray image of the left ventricle of the heart. The largest structures inside the circular area representing the borders of the image intensifier are: the left ventricle with the descending aorta (centre), the diafragma muscle (lower left) and a metallic filter (upper right). To perform the segmentation of this kind of images we utilised three layers of neurons: one to represent the image intensifier, one for the background (soft tissues with a low density) and one for the structures just mentioned (they have approximately the same value of D_c). Figure G5.5.4 (right) shows the activation of this last layer. Diafragma muscle, left ventricle with aorta and metallic filter have been correctly represented as disjunct regions. Another example is given in Figure G5.5.5 which illustrates the segmentation of a radiogram of a finger. Although, in this case, the network




Figure G5.5.2. Segmentation of an MR image of the thorax.




Figure G5.5.3. Multiscale segmentation of an MR image of the head.

has not been capable of splitting the bone part of the finger into its anatomical components because of the very limited inter-bone space, the important discrimination between soft tissue and bone is correct, even where bone and soft tissue overlap.




Figure G5.5.4. Segmentation of a cine-angiographic image of the left ventricle.



Figure G5.5.5. Segmentation of a radiogram of a tract of a finger.

G5.5.6 Conclusion

In this chapter we have described a neural architecture for the segmentation of medical images. With simple topology and parameter changes the architecture can be adapted to perform the 2-D, 3-D and multi-scale segmentation of tomographic and X-ray images. Thanks to its broad applicability, to the robustness and sensitivity shown in the experiments and to its implementability with fine-grained parallel hardware, this architecture seems to meet the requirements to be considered a possible general purpose solution to the problem of medical image segmentation.

Acknowledgements

This work has been partially supported by Italian Ministry for University and Scientific and Technologic Research (MURST).

References

- S.C. Amatur, D. Piraino, and Y. Takefuji. Optimization neural networks for the segmentation of magnetic resonance images. *IEEE Transactions on Medical Imaging*, 11(2), 1992.
- G.L. Bilbro, M. White, and W. Snyder. Image segmentation with neurocomputers. In R. Eckmiller and C. v.d. Malsburg, editors, *Neural Computers*, Berlin, 1987. Springer-Verlag.
- J.Y. Catros and D. Mischeler. An artificial intelligence approach for medical picture analysis. *Pattern Recognition Letters*, 8:123–130, 1988.
- C. T. Chen, E. C. K. Tsao, and W. C. Lin. Medical image segmentation by a constraint satisfaction neural network. *IEEE Transaction on Nuclear Science*, 38(2):678–686, 1991.
- G. Coppini, M. Demi, R. Poli, and G. Valli. An artificial vision system for X-ray images of human coronary trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(2):156–162, 1993.
- G. Coppini, R. Poli, M. Rucci, and G. Valli. A neural network architecture for understanding 3D scenes in medical imaging. *Computer and Biomedical Research*, 25:569–585, 1992.
- T. Darrell, S. Sclaroff, and A. Pentland. Segmentation by minimal description. In *IEEE International Conference on Computer Vision III*, Osaka, Japan, 1990.
- R. Deklerck, J. Cornelis, and M. Bister. Segmentation of medical images. *Image and Vision Computing*, 11(8):486–503, October 1993.
- G. Gerig, J. Martin, R. Kikinis, O. Kubler, M. Shenton, and F. A. Jolesz. Unsupervised tissue type segmentation of 3D dual-echo MR head data. *Image and Vision Computing*, 10(6):349–360, 1992.
- W. E. Higgins, N. Chung, and E. L. Ritman. Extraction of left-ventricular chamber from 3-D CT images of the hart. *IEEE Transactions on Medical Imaging*, 9(4):384–395, 1990.
- J.J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81:3088–3092, 1984.
- J.J. Hopfield and D.W. Tank. “Neural” computation of decisions in optimization problems. *Biological Cybernetics*, 52:141–152, 1985.
- J.J. Hopfield and D.W. Tank. Computing with neural circuits: a model. *Science*, 233:625–633, 1986.

-
- M. Joliot and B. M. Mazoyer. Three-dimensional segmentation and interpolation of magnetic resonance brain image. *IEEE Transactions on Medical Imaging*, 12(2):269–277, June 1993.
- JR. J. W. Klingler, C. L. Vaughan, JR. T. D. Franker, and L. T. Andrews. Segmentation of echocardiographic images using mathematical morphology. *IEEE Transactions on Biomedical Engineering*, 35(11):925–935, 1988.
- T. Lei and W. Sewchand. Statistical approach to X-ray CT imaging and its applications in image analysis—Part II: A new stochastic model-based image segmentation technique for X-ray CT image. *IEEE Transactions on Medical Imaging*, 11(1):62–69, 1992.
- C. Li, D. B. Goldgof, and L. O. Hall. Knowledge-based classification and tissue labeling of MR images of human brain. *IEEE Transactions on Medical Imaging*, 12(4):740–750, December 93.
- G. Manos, A. Y. Cairns, I. W. Ricketts, and D. Sinclair. Automatic segmentation of hand-wrist radiographs. *Image and Vision Computing*, 11(2):100–111, March 1993.
- D. Marr. *Vision*. W.H. Freeman & Co., New York, 1982.
- M. Özkan, B. M. Dawant, and R. J. Maciunas. Neural-network-based segmentation of multi-modal medical images: A comparative and prospective study. *IEEE Transactions on Medical Imaging*, 12(3):534–544, September 1993.
- R. Poli, G. Coppini, and G. Valli. Recovery of 3-D closed surfaces from sparse data. *Computer Vision Graphics and Image Processing: Image Understanding*, 60(1):1–25, July 1994.
- S. V. Raman, S. Sakar, and K. L. Boyer. Hypothesizing structures in edge-focused cerebral magnetic resonance images using graph-theoretic cycle enumeration. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 57(1):81–98, January 1993.
- S. P. Raya. Low-level segmentation of 3-D magnetic resonance brain images — a rule-based system. *IEEE Transactions on Medical Imaging*, 9(3):327–337, 1990.
- T. R. Reed. Region growing using neural networks. In Harry Wechsler, editor, *Neural Networks for Perception, Vol. 1*, pages 386–397. Academic Press, San Diego, CA, 1992.
- S. R. Reuman and D. D. Hoffman. Regularities of nature: the interpretation of visual motion. In Alex P. Pentland, editor, *From Pixels to Predicates*, pages 201–226. Ablex, Norwood, New Jersey, 1986.
- R. H. Silverman and A. S. Noetzel. Image processing and pattern recognition in ultrasonograms by backpropagation. *Neural Networks*, 3:593–603, 1990.
- J. G. Thomas, R. A. Peters II, and P. Jeanty. Automatic segmentation of ultrasound images using morphological operators. *IEEE Transactions on Medical Imaging*, 10(2):180–186, June 1991.
- D. L. Toulson and J. F. Boyce. Segmentation of MR images using neural nets. *Image and Vision Computing*, 10(5):324–328, June 1992.
- T. Wang, X. Zhuang, and X. Xing. Robust segmentation of noisy images using a neural network model. *Image and Vision Computing*, 10(4):233–240, 1992.