

Neural Inhabitants of MR and Echo Images Segment Cardiac Structures

Riccardo Poli and Guido Valli

Department of Electronic Engineering
University of Florence (Italy)

Abstract

This paper describes a new approach to the problem of the segmentation of cardiac structures in medical imaging. The approach is based on the idea of breeding and selecting artificial creatures who live in such images and are fed with the boundaries of the structures to be segmented. Our creatures, the Gnets, are simple individuals based on recurrent neural networks who can see, know their position in the environment, move inside it and eat. Their behavior is developed through a genetic algorithm which keeps a population of Gnets and mates the best individuals. Performance is evaluated on a set of test images of known segmentation. Preliminary results of this approach are reported.

1 Introduction

Segmentation of cardiac chambers is the first step for almost any kind of automatic high-level analysis of heart shape and function. Unfortunately, segmenting medical images is a very difficult task. This is due to noise, masking structures, biological shape variability, tissue inhomogeneity, imaging-chain anisotropy and variability, etc. In order to overcome these problems most researchers have adopted the strategy of exploiting anatomical knowledge about the imaged structures.

Artificial Neural Networks (ANNs) with supervised training [1] have revealed particularly suited to implement this strategy [2, 3]. Their advantages are robustness (derived from the parallel and distributed nature of the processing they perform) and flexibility (coming from the capability of learning by examples). However, three problems presently reduce the practical applicability of this kind of ANN: a) the remarkable computation load involved in a pixel-by-pixel ANN-based segmentation, b) the complexity of setting up a training set containing the correct teaching

input for each output neuron, and c) the constraint of using feedforward connection topologies. While a) can be overcome with special-purpose hardware or hybrid architectures, b) and c) can seriously limit the kind of segmentation strategies which can realistically be implemented with current ANNs.

In this paper we present a new approach to cardiac image segmentation which can overcome the aforementioned drawbacks. It is based on the idea, borrowed from the field of Artificial Life [4], of selecting artificial creatures by means of genetic algorithms on the basis of their fitness to the surrounding environment [5, 6, 7]. Our creatures, we call them Gnets (Genetic Networks), are basically a kind of recurrent ANNs which “lives” inside the image to be segmented (the environment). Actually, Gnets can recognize the boundaries of the structures to be segmented (through a retina), are aware of their position in the environment, move inside it and open their “mouth” to eat. Their behavior is developed through a genetic algorithm [8, 9] which keeps a population of Gnets and mates the individuals which show the best performance on a set of test images of known segmentation.

The architecture of Gnets, the procedure for their breeding, and preliminary results of this approach are described in the next sections.

2 Gnets' Architecture

Gnets are recurrent ANNs, the architecture of which Figure 1 shows of a simplified version.

They include two kinds of sensory (input) neurons: visual neurons belonging to the retina of an “eye” through which Gnets “see” the surrounding environment and proprioceptive neurons through which Gnets know their position inside the environment. The neurons of the retina are organized in a 5×5 layer and receive, as input, the gray levels of the pixels lying in a 5×5 neighborhood of the current position




Figure 1: Architecture of Gnets

of the Gnet in the image. Proprioceptive neurons are organized in two sets of 8 units, which encode the x and y coordinates of the Gnet in the environment. The encoding is as follows. First each coordinate c is normalized in the range $[0, 7]$: for a $N_p \times N_p$ environment, c is transformed into $c_n = c \times 8/N_p$. Then the neurons of index $\lfloor c_n \rfloor$ and $\lfloor c_n \rfloor + 1$, $\lfloor \cdot \rfloor$ being the integer-part operator, are activated proportionally to $1 - (c_n - \lfloor c_n \rfloor)$ and $c_n - \lfloor c_n \rfloor$, respectively.

The input neurons are connected with forward links to a set of hidden neurons which, along with five motor neurons, make up the “brain” of each Gnet. Inside the brain there are both forward and backward connections which provide Gnets with memory: the state of the brain at a given time depends on the current sensory input and on the previous states (i.e. on the Gnet’s past life).

The five motor neurons have the following functions. Four of them control the motion of the Gnets inside the environment: they represent the actions “go up”, “go down”, “go left” and “go right”. Note that no limit is imposed on Gnets motion as the environment has a toroidal shape (when a Gnet gets out of the environment, it reenters the opposite side). The fifth motor neuron control the opening and closing of Gnets’ mouth.

3 Breeding Gnets

In order to obtain Gnets capable of segmenting cardiac structures we have to find out a set of weights and initial activations which provide them with the correct behavior. The ideal behavior would be first searching the image for the structure of interest with the mouth closed and then moving along its boundaries with the mouth opened (eating). Unfortunately, this problem does not allow for either a mathematical formulation (which would allow using analytic optimization techniques) or a conventional training set (whose compilation, if possible, would be actually too much time consuming). Therefore, we have adopted an optimization algorithm which does not require any of them: a genetic algorithm.

Genetic algorithms (GAs) are optimization procedures inspired by the natural process of selection and by genetics [8]. Virtually, all GAs operate as follows: a) randomly generate a set (population) of solutions (individuals) for the problem, b) evaluate the goodness (fitness) of each individual, c) generate a new population partly by cloning the fittest individuals, partly by combining and partly by mutating them, and d) iterate steps b–d until a stop condition is true.

The most common type of GA requires each solution of the problem (phenotype) to be encoded as a bit string (genotype). In this case, step b) involves decoding and evaluating each individual, which are usually the computationally most expensive parts of GAs. In general, when building a new generation (step c), the algorithm copies part of the strings of the previous population and chooses pairs of them (parents) for generating new solutions (offsprings). This is performed through a genetic operation called *crossover*. As shown in Figure 2a, crossover involves cutting two bit strings at a (randomly chosen) common point and exchanging the left (or right) hand side sub-strings. This search operator, along with the fact that individuals are selected for cloning and crossing-over with a probability proportional to their fitness, provide new generations containing better and better individuals. Another genetic operator, called *mutation*, is also applied (though with a very low probability) on the individuals of a new generation. Mutation can improve the search for solutions by reverting random bits of the genotype (see Figure 2b).

In order to make a genetic algorithm force Gnets to behave correctly, we have to define a proper fitness function f . This is a crucial step for making Gnets work. We have designed a biologically inspired fitness function. Basically, f can be considered a sort of “vital energy” given to each Gnets. Primarily, it is determined by the “food” each Gnet has been capable

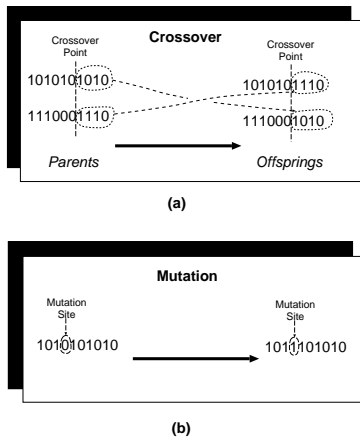


Figure 2: Genetic operators.

of eating during his life. Food is represented by the pixels representing the borders of the structure to be segmented. In order to prevent Gnets from evolving the “optimum” strategy of going around with their mouths opened, we punish them if they eat “ground” (pixels not belonging to the structure of interest). Of course, eaten food is removed from the environment to prevent Gnets from stopping where they have found it. Finally, f is decreased by each movement the Gnet perform. Gnets’ life starts, with a fixed amount of vital energy f , from a random position in the image. Life lasts for a fixed number N_t of time steps unless $f = 0$. In such a case the Gnet dies prematurely. In order to make the evaluation of the fitness of each Gnet independent of the starting position, fitness is evaluated as the mean on N_t lives. For each life, the environment is randomly selected from a set of images of the same kind so as to make the Gnet independent of the changes of shape and contrast of the structures to be segmented.

The evaluation of the fitness function just described requires the exact position of the food contained in the environment to be known. This means that, during the genetic selection-phase, images of known segmentation (training set) are to be used as environment. However, Gnets do not need any pixel-by-pixel teaching input, as they do not have any sensory neuron revealing the actual presence of food in a given pixel or in their “stomach”. So, after the genetic algorithm has developed a set of Gnets with the proper behavior, segmentation can simply be performed by tracing the moves of those Gnets inside the image to be segmented and labeling the pixels where one or more of them have opened their mouths.

4 Preliminary Results

Genetic algorithms have a broad applicability, but, unfortunately, they are not very efficient. To make things worse, in Artificial-Life experiments, evaluating the fitness of an individual involves making him live for hundreds or thousands of time steps. In addition, such an evaluation is performed for the hundreds of individuals of a population, and for the hundreds of generations needed to develop any non-trivial behavior. The huge computation power required by such procedures is presently hardly available on supercomputers and fine-grained parallel computers which are reported to carry out the task in weeks or months of CPU time [5, 6, 7]. Our experiments with Gnets are being carried out by two IMB RISC 6000 workstations, which, though quite powerful, provide us with 1–2 orders of magnitude less performance than the above-mentioned computers.

We are experimenting Gnets on MR images of the thorax and echo images of the heart. In both cases Gnets have 41 sensory units, 18 hidden units and 5 motor neurons. The weights associated with synaptic connections, the starting states and the biases of the 23 neurons included in Gnets’ brain are the parameters which undergo genetic optimization. Weights and biases are fixed point numbers in the range $[-10, +10]$; activations are in the interval $[0, 1]$. Each parameter is encoded as a 6 bit Gray-coded integer. As a result, Gnets’ chromosomes contain about 4,000 bits. Fitness is evaluated on $N_t = 5$ lives lasting for $N_t = 256$ time steps each. Populations include 300 Gnets, and are bred for a maximum of 1,000 generations. Given the limited number of individuals in each population, there is no certainty of obtaining Gnets behaving correctly in any single experiment. Therefore, we adopt the strategy of going on repeating the experiments until a good behavior is developed.

Figure 3 shows the plot of the fitness of a population of Gnets inhabiting an echo image.

The path followed by a Gnet is superimposed (in light gray) on the binary image; the pixels in which the Gnet has opened his mouth are drawn in black. While it is clear that this Gnet alone is able to segment only part of the LV boundary, other Gnets show complementary behaviors, so that the complete segmentation can be obtained by a “team” of Gnets carefully selected.

Figure 4 shows a Gnet hunting inside a two-chamber echocardiographic image. This is a harder task for Gnets, as no preliminary binarization is performed. Images are decimated with smoothing to a resolution of 64×64 pixels, and Gnets are fed with the atrial and ventricular boundaries. Again




Figure 3: Gnets' fitness vs. generation.




Figure 4: The behavior of a Gnet living inside an echocardiographic image.

only part of the structures of interest has been segmented. Unfortunately, some false detections have also been performed. However, in this case too, there are Gnets who like other parts of the structures to be segmented, so that a team strategy can be pursued. In addition, off-sector detections can be ignored.

5 Conclusions

This paper describes a new approach to the segmentation of cardiac structures in medical images. It is based on artificial, (moderately) intelligent creatures which efficiently scan the image looking for the appetizing boundaries of heart chambers. They are developed through the selective pressure of an evolutionary strategy.

Though we are aware that the results reported in this paper are preliminary, the continuous improvement of the hunting capability of Gnets and their considerable efficiency suggest that our approach could be a right answer to the problems encountered in automating the analysis of medical images.

Acknowledgements

This work has been partially supported by funds of the Italian MURST and CNR.

References

- [1] D. Rumelhart and J. McClelland, eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1-2*. Cambridge, MA: MIT Press, 1986.
- [2] R. Poli, G. Coppini, R. Nobili, and G. Valli, "Lv shape recovery from echocardiographic images by means of computer vision techniques and neural networks," in *Computers in Cardiology*, (Venice), pp. 117–120, IEEE Computer Society press, 1991.
- [3] G. Coppini, R. Poli, M. Rucci, and G. Valli, "A neural network architecture for understanding 3d scenes in medical imaging," *Computer and Biomedical Research*, vol. 25, pp. 569–585, 1992.
- [4] C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds., *Artificial Life II*. Reading, Massachusetts: Addison-Wesley, 1992.
- [5] B. Fullmer and R. Miikkulainen, "Using marker-based genetic encoding of neural networks to evolve finite-state behaviour," in *Proceedings of the First European Conference on Artificial Life (ECAL-91)*, 1991.
- [6] R. J. Collins and D. R. Jefferson, "An artificial neural network representation for artificial organisms," in *Parallel Problem Solving from Nature* (R. Männer and D. E. Goldberg, eds.), Springer-Verlag, 1991.
- [7] R. J. Collins and D. R. Jefferson, "Antfarm: Towards simulated evolution," in *Artificial Life II* (C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds.), pp. 579–601, Addison-Wesley, 1991.
- [8] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, Massachusetts: Addison-Wesley, 1989.
- [9] T. Bäck and H.-P. Schwefel, "An overview of evolutionary algorithms for parameter optimization," *Evolutionary Computation*, vol. 1, no. 1, 1993.

Address for correspondence:

Riccardo Poli, PhD

Department of Electronic Engineering

University of Florence

Via S. Marta, 3

I-50139 – Florence

Italy

E-mail address: poli@viper.ing.unifi.it