ESANN'2004 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 28-30 April 2004, d-side publi., ISBN 2-930307-04-8, pp. 269-274

Neural network-based calibration of positron emission tomograph detector modules

B. Lazzerini, F. Marcelloni, G. Marola, S. Galigani

Dipartimento di Ingegneria dell'Informazione, Via Diotisalvi 2, 56122 Pisa, Italy

Abstract. In this paper we describe a neural network-based method aimed at automatically calibrating the detector module contained in a scanner for a high-resolution positron emission tomography (PET) system for small animals. The detector module is composed of crystal elements, arranged in a regular matrix and sensitive to gamma rays emitted by a radioactive source. The crystal matrix is optically coupled to a position-sensitive photo-multiplier tube, which reconstructs the original image. Calibration, required to cope with spatial distortions introduced by the optical system, consists of a segmentation process of the image produced after the photo-multiplier tube into a fixed number of areas. The purpose of this segmentation is to map each pixel of the perceived image onto the pertinent crystal, which was actually struck by the gamma ray emitted by the radioactive source.

1. Introduction

Positron emission tomography (PET) is a nuclear medicine technique which allows localizing with high precision, within the brain or other body organs, a substance, previously injected into the patient's body and marked with a radioisotope emitting positrons [1]. When positrons collide with electrons, gamma ray photons are produced. More precisely, each nuclear reaction gives origin to two gamma rays which are paired so as to move away in exactly opposite directions. A (frequently low) percentage of these pairs of gamma rays are sensed simultaneously by two detectors, which, being placed on opposite sides, can intercept the flight line of such pair. When hit by gamma rays, the two detectors, usually based on scintillator crystals, emit light photons which are captured and amplified by a position-sensitive photo-multiplier tube (PSPMT). Trough an appropriate decoding circuit, the photo-multiplier establishes the coordinates of the source of each light photon. In this way, variations in radiation intensity and in spatial location at point sources in the body activate the detector module that maps the radiation intensity in the X-Y space to create an image.

PET scanners allow following the path of the radioisotope throughout the body and to examine the organ where a particular radioisotope may concentrate preferentially. Of course, high sensitivity is necessary both to image radiolabelled substances, which are distributed in live subjects in very low quantities, and to image fast dynamic processes *in vivo*. A number of different approaches to achieving high sensitivity have been employed, including the development of novel scintillator crystals. In particular, YAP:Ce (yttrium aluminium perovskite activated by cerium) scintillators, providing good resolution and sensitivity, and having very low energy secondary X-ray emissions, are particularly suitable for this kind of imagining applications.

ESANN'2004 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 28-30 April 2004, d-side publi., ISBN 2-930307-04-8, pp. 269-274

PET scanners are more and more frequently used in neurology, cardiology and tumour detection in various parts of the human body. High-resolution PET scanners have also been developed for imaging small laboratory animals, such as rats and mice, as an important tool in biomedical research for drug discovery and development, and evaluation of therapeutic efficacy in small animal models of human disease [2].
In this paper, we refer to a positron emission tomograph for small animals. The tomograph includes a four-head detector module consisting of four matrices of 400 YAP:Ce crystals optically coupled to as many PSPMTs.
One of the problems encountered when using PET systems, especially in scanners employed in experiments on small animals, comes from various kinds of distortions of the final image mainly due to irregularities of the optical-electronic system. In fact, an appropriate calibration of the resulting image is necessary.
Essentially, calibration aims to associate each pixel of the image produced by a detector module with the crystal influencing that pixel. The calibration process is operated on the image obtained by illuminating each detector module with a uniform and parallel beam of gamma rays for a period of time sufficient to allow high-resolution image acquisition. Being the intensity level of the produced image proportional to the number of gamma rays that hit the corresponding crystal, the identification of the region of influence of each crystal can be guided by the search of intensity peaks in the image. The calibration process is quite often performed through direct intervention by the user resulting in a highly time-consuming operation. Otherwise, procedures for automatic calibration have also been developed (e.g., a calibration algorithm based on the watershed transform is presented in [3]), but they usually rely on the interactive correction of the results by the user.
In this paper, we present a neural network-based method for automatic calibration of PET detector modules without requiring any user's intervention. The method exploits a self-organizing map (SOM) [4] which retains the main characteristics of the SOM standard model though introducing a new concept of neighbourhood and a further, purposely-defined learning phase. In the following, we describe the key features of the method and present experimental results.

2. The method

In self-organizing maps, neurons learn to recognize both the distribution and the topology of the input vectors they are trained on. The neurons are arranged originally in physical positions in the one- or two-dimensional map according to a specific topology that defines which neurons are in fact neighbours of the winning neuron. The neighbourhood function may take different forms, such as squared or hexagonal. During learning, both the winning neuron and all neurons within a specified neighbourhood are updated by means of the Kohonen learning rule, which moves the weights of such neurons closer to the input vector. The weights are altered proportional to the learning rate. Both the learning rate and the neighbourhood size are appropriately decreased during training. Typically, the neighbourhood is selected fairly wide in the beginning and then permitted to shrink with time. Similarly, the learning rate parameter starts with a value close to unity and then decreases gradually. Two successive distinct learning phases can be distinguished: the *ordering phase*,

during which the topological ordering of the weight vectors takes place, and the *tuning* (or *convergence*) *phase*, during which fine tuning of the computational map is achieved by maintaining a fixed small neighbourhood and a slowly decreasing learning rate, without altering the ordering learned in the preceding phase.
From the previous simple review of the main characteristics of self-organizing maps, it turns out that we need to make appropriate changes to the standard model of self-organizing map in order to reach our purpose. Indeed, each detector module of the PET system we refer to consists of a squared matrix made of 20x20 crystals sensitive to gamma rays. Owing to measurement noise, a not perfectly linear behaviour of the crystals, possible border effects, and non-uniform sensitivity of the multiplier tubes, the resulting image looks like a geometrically distorted squared grid. Therefore, we basically aim to rebuild the grid structure of the detector crystal matrix by identifying cluster centres and linking these centres in orderly way according to a fixed connection topology. Neurons in the final map must be spaced out in a fairly uniform way consistently with the geometric distortions present in the produced image and independently of the difference in the intensity levels associated with the peaks.
To this aim, we introduce an intermediate phase, called *grid ordering phase*, between the standard ordering and tuning phases. During this phase, the learning procedure tries to allocate the neuron grid on the input vectors by heavily exploiting the information pertinent to the expected grid structure. The basic idea is to define a more specific concept of neighbourhood which takes into account the geometric nature of the neural grid, and to appropriately update the weights of the winner and the neurons inside this specific neighbourhood. Generally, the neighbourhood function is taken to include neighbours in a (typically squared) region, called *activity bubble*, around the winning neuron. During the grid ordering phase, the neighbourhood of the winner consists of the neurons that are within a certain distance from the winner and at the same time on either the same row (coordinate X) or the same column (coordinate Y) of the neural grid as the winner. This choice aims to assign a high weight to the grid topology originally associated with the centres of the intensity peaks.

3. The learning process

The first learning phase is the ordering phase, which maps the weight vectors so as to loosely represent the distribution surface in the input space: the aim is to approximate the shape of the image, previously compensated and filtered. The weight map, initially characterized by uniform distribution, is warped owing to the neuron attraction by the pixels presented to the network. The learning rate varies linearly between two user-specified values; the same holds for the neighbourhood distance.
The following learning phase is the grid ordering phase, which stems from a mechanical view of the clustering problem. Assuming that a fixed-mass small ball is associated with each neuron, each ball is linked to the adjacent balls through a spring system along the two dimensions X and Y. Each ball, except for those on the edges, has four immediate neighbouring balls. The problem of finding the intensity peak centres can be transposed into the complementary approach of interpreting high intensity levels as valleys: a pixel with a high intensity value corresponds to a deep point in the valley. Our aim is therefore to make the 400 balls fall into the best 400

ESANN'2004 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 28-30 April 2004, d-side publi., ISBN 2-930307-04-8, pp. 269-274

valleys detectable in the input surface.
Of course, the fact that springs are present only on the X and Y directions, and not on diagonal trajectories, stresses the importance of the grid morphology of the expected result. When a ball (actually, a neuron) is attracted by a pixel relative to a surface depression (actually, a cluster), this ball not only tries to occupy the valley central position but, through the spring system, applies an appropriate force to the adjacent balls. More precisely, let *Dx* and *Dy* be, respectively, the horizontal and vertical displacements of the winning ball to get closer to the training pixel. All the balls in the winner's current neighbourhood are moved according to the following rule: each ball on the same row as the winner is moved in the same direction as the winner along the horizontal line, this shift is proportional to *Dx* (called *primary horizontal factor*) and inversely proportional to the distance between the ball and the winner; the same occurs for the vertical direction. The *primary* movement, which is either horizontal or vertical depending on the fact that the neighbouring ball is on the same row or column as the winner, is accompanied by a *secondary* movement along the orthogonal direction; this secondary movement is originated from the grid settling following the movement of the winner in both the horizontal and vertical directions. The secondary movement, made by a ball along the direction orthogonal to that of its primary movement, is proportional to a *secondary* (*horizontal* or *vertical*) *factor*, which is smaller than the corresponding primary factor, and inversely proportional to the distance between the ball and the winner. A substantial difference between primary and secondary movements is the shift direction: whereas a primary movement causes a shift of all the neighbouring balls in the same direction as that of the winner along that dimension, a secondary movement represents the attractive force by the training pixel and consequently the neighbouring balls move towards the winner. In this way, e.g., a horizontal spring has a primary effect along the horizontal direction and a secondary effect along the vertical direction.
In the grid ordering phase, the weights of the winner are updated proportional to the current value of the learning rate while those of the neighbouring neurons are updated by summing the primary and secondary movements still proportional to the current learning rate. During the search for the best distribution it may happen that two balls are so close that they are attracted within the same valley. This has a negative effect on the cluster identification. We can avoid this problem by exploiting information about the average distances between adjacent balls. More precisely, we periodically calculate the average distances along the horizontal and vertical directions and check whether the distance between any two adjacent balls falls below a threshold which is a function of the average distances themselves. When this occurs, the two balls undergo an elastic collision which moves them away in opposite directions together with all the balls arranged along the direction involved in the collision. The information about the average distance between adjacent balls are also used to reveal situations in which two adjacent balls move away too far from each other along one direction. In this case, all the balls on the detected direction undergo an appropriate attractive force.
The final learning phase is the tuning phase during which the neighbourhood of each neuron consists of the neuron itself: in practice, each input vector causes the updating of the winner's weights only. On the other hand, the grid ordering phase has placed each neuron in the best influence zone with respect to the crystal matrix and the task

ESANN'2004 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 28-30 April 2004, d-side publi., ISBN 2-930307-04-8, pp. 269-274

of the tuning phase is just a fine tuning of the search for the best centre to associate with the intensity peak present in that zone.

4. Experiments and results

We use a network with 400 neurons arranged in the form of a two-dimensional lattice with 20 rows and 20 columns. The network is trained with a two-dimensional input vector, whose elements represent the original image pixels normalized within the range [-1,1]. The initial synaptic weights are uniformly distributed within the two-dimensional input space. The neighbourhood distance is the Manhattan distance.
The input data consist of a 400*400 pixel image produced by the photo-multiplier tube. For each pixel, the intensity of the energy level is provided, which is proportional to the number of photons revealed during the exposition time to gamma rays. We carried out several experiments with both good-quality and low-quality images. In the former case, the image consists of well distinguished high-intensity peaks, whereas in the latter case, the image contains various kinds of distortions such as missing or fused peaks. An example of low-quality image is displayed in figure 1(a), where brighter zones correspond to higher intensity levels. We can notice a wide range of intensity levels; further, not all peaks are well separated, and some of them, especially those close to the image edges, can be scarcely detected. In figure 1(b) the neural grid produced at the end of the ordering phase is placed on the original uncompensated and unfiltered image. The neurons are represented as white spots. Figure 1(c) shows the grid layout after the grid ordering phase. Finally, Figure 1(d) shows the grid resulting from the tuning phase. We can notice that, in spite of the distortions present in the original image, the self-organizing map is able to keep the lattice topology that characterizes the peak distribution. The results are excellent not only in identifying the internal peaks, but also those located at the edges of the image. The system was developed in Matlab 6.5. The three learning phases last 2, 120 and 20 seconds, respectively, on average on an Intel Pentium 4, 3 GHz PC.

5. Conclusions

We have presented a SOM-based system for image calibration in a PET scanner used for experiments on small animals. Essentially, we build a mechanical model of the calibration system consisting of balls, associated with the neural network neurons, and springs arranged in a grid structure. The problem of identifying the image intensity peaks is transformed into a search for a ball distribution to 400 valleys, representing pixel clusters, on the surface made of the image pixels. During training, the balls, stimulated by the input pixels, start vibrating until a stable equilibrium point is reached. The final equilibrium corresponds to have each ball in the appropriate valley. Experimental results have shown high accuracy and effectiveness in identifying the crystal influencing each image pixel, low calibration time, and robustness against noise, spatial distortions and large variability in the intensity degrees of the pixels. The segmentation process is highly reliable both in the internal regions and on the borders of the images. Finally, no interactive correction by the user is required.

ESANN'2004 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 28-30 April 2004, d-side publi., ISBN 2-930307-04-8, pp. 269-274

Acknowledgment

The images for the experiments were provided by I.S.E. S.r.l., Vecchiano (Pisa) Italy. They were produced by the YAP-(S)PET scanner manufactured by I.S.E.

References

1. J.L. Humm, A. Rosenfeld, A. Del Guerra: From PET detectors to PET scanners. European J. Nuclear Medicine and Molecular Imaging, 30, 1574-1597 (2003)
2. M.E. Phelps: PET: the merging of biology and imaging into molecular imaging. Journal Nuclear Medicine, 41, 661-681 (2000)
3. A. Mao: Positron emission tomograph detector module calibration through morphological algorithms and interactive correction. 1999-2000 International Student Technical Writing Competition (ISTWC) winners
4. S. Haykin: Neural Networks: A Comprehensive Foundation. New York, Macmillan (1994)

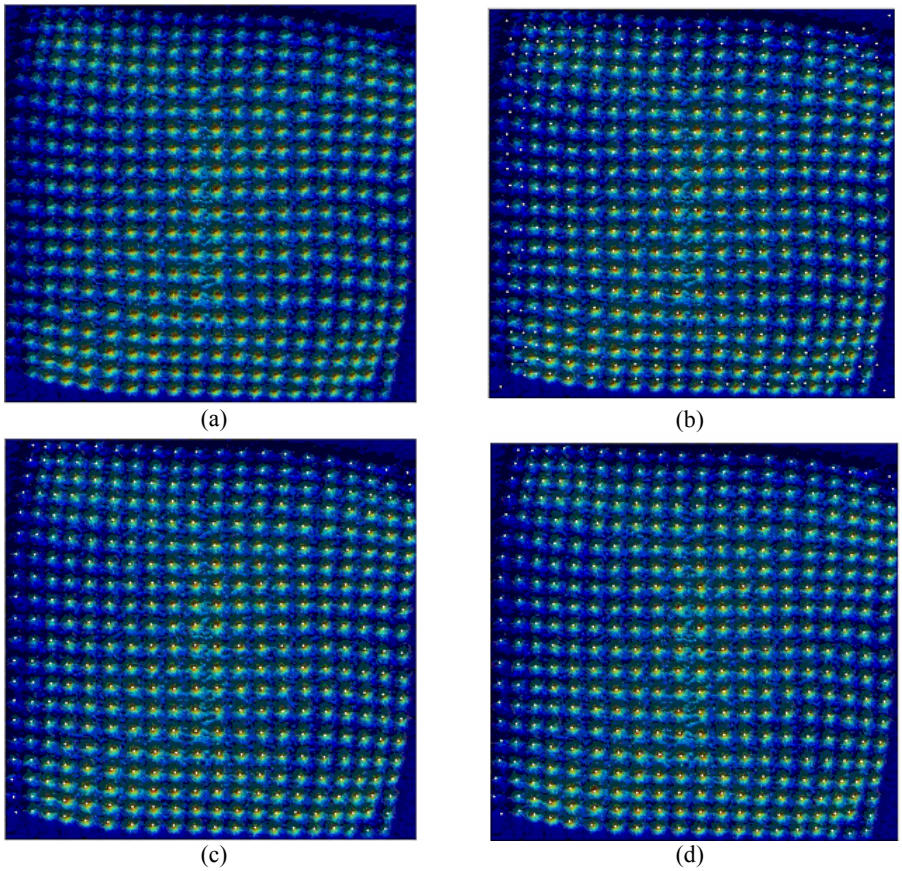

Fig. 1. (a) Original image; (b) Neural grid (after the ordering phase) superimposed on the original image; (c) Neural grid (after the grid ordering phase) superimposed on the original image; (d) Neural grid (after the tuning phase) superimposed on the original image.