A Neural Network that Learns to Interpret Myocardial Planar Thallium Scintigrams*

Charles Rosenberg, Ph.D.†
Department of Computer Science
Hebrew University
Jerusalem, Israel
`crr@cogsci.psych.utah.edu`

Jacob Erel, M.D.
Department of Cardiology
Sapir Medical Center - Meir General Hospital
Kfar Saba, Israel

Henri Atlan, M.D., PhD.
Department of Biophysics
and Nuclear Medicine
Hadassah Medical Center
Jerusalem, Israel

June, 1992

Abstract

The planar thallium-201 (^{201}Tl) myocardial perfusion scintigram is a widely used diagnostic technique for detecting and estimating the risk of coronary artery disease. Interpretation is currently based on visual scoring of myocardial defects combined with image quantitation and is known to have a significant subjective component.

Neural networks learned to interpret thallium scintigrams as determined by both individual and multiple (consensus) expert ratings. Four different types of networks were explored: single-layer, two-layer back-propagation (BP), BP with weight smoothing, and two-layer radial basis function (RBF). The RBF network was found to yield the best performance (94.8% generalization by region) and compares favorably with human experts.

We conclude that this network is a valuable clinical tool that can be used as a reference "diagnostic support system" to help reduce inter- and intra-observer variability. This system is now being further developed to include other variables that are expected to improve the final clinical diagnosis.

1 Introduction

Coronary artery disease (CAD) is one of the leading causes of death in the Western World. The planar thallium-201 is considered to be a reliable diagnostic tool in the detection of CAD. Thallium is a radioactive isotope that distributes in mammalian tissues after intervenous administration and is imaged by a gamma camera. The resulting scintigram is visually interpreted by the physician for the presence or absence of defects — areas with relatively lower perfusion levels. In myocardial applications, thallium is used to measure myocardial ischemia and to differentiate between viable and non-viable (infarcted) heart muscle [16].

Diagnosis of CAD is based on the comparison of two sets of images, one set acquired immediately after a standard effort test (BRUCE protocol), and the second following a delay period of four hours. During this delay, the thallium redistributes in the heart muscle and spontaneously decays. Defects caused by scar tissue are relatively unchanged

*This paper is to appear in *Neural Computation*.

†Supported by grants from the Ministry of Science and Technology, Israel and from the Golda Meir Fund. Current address: GRECC 182, VA Medical Center, 500 Foothill Blvd., Salt Lake City, Utah 84148

over the delay period (fixed defect), while those caused by ischemia are partially or completely filled-in (reversible defect) [2].

Image interpretation is difficult for a number of reasons: the inherent variability in biological systems which makes each case essentially unique, the vast amount of irrelevant and noisy information in an image, and the "context-dependency" of the interpretation on data from many other tests and clinical history. Interpretation can also be significantly affected by attentional shifts, perceptual abilities, and mental state [10, 6].

Neural networks have been applied to several problems in cardiology including the prediction of stenosis [17, 5, 4, 12]. These studies encouraged us to explore the problem of visual interpretation using neural networks which, to our knowledge, has not been previously addressed.

2 Data

Scintigraphic images were acquired for each of three views: anterior (ANT), left lateral oblique (LAO 45), and left lateral (LAT)[1] for each patient case. Each image was first pre-processed[2] and presented to the networks as circumferential profiles [9, 7][3], in which maximum pixel counts within each of 60, 6° contiguous segmental regions are plotted as a function of angle [8] (see Figure 1 A).

Cases were pre-selected based on the following criteria [2]:

- **Insufficient exercise.** Cases in which the heart rate was less than 130 b.p.m. were eliminated, as this level of stress is generally deemed insufficient to accurately distinguish normal from abnormal conditions.
- **Positional abnormalities.** In a few cases, the "region of interest" was not positioned or aligned correctly by the technician.
- **Increased lung uptake.** Typically in cases of multi-vessel disease, a significant proportion of the perfusion occurs in the lungs as well as in the heart, making it more difficult to determine the condition of the heart due to the partially overlapping positions of the heart and lungs.

Twenty percent of the cases were eliminated due to insufficient heart rate, and an additional 5 percent due to either positional abnormalities or increased lung uptake. A set of one hundred usable cases resulted.

Each case was visual scored for each of nine anatomical regions generally accepted as those that best relate to the coronary circulation: *Septal: proximal* and *distal*, *Anterior: proximal* and *distal*, *Apex*, *Inferior: proximal* and *distal*, and *Posterior-Lateral: proximal* and *distal*. Scoring for each region was from *normal* (1) to *severe* (4), indicating the level of the observed perfusion deficit.

Approximately 80% of the cases were visually interpreted by a single expert alone. Ten percent were read by all three experts, and 10% by two experts, with agreement in the scoring reached by consensus.

2.1 Methods

Four different types of networks were explored: a one-layer (sigmoidal) network trained using the delta rule, a two-layer network trained using back-propagation (BP), a two-layer network trained using BP combined with weight smoothing (to be described), and a two-layer network combining one layer of competitive radial-basis function (RBF) units and a second layer of (sigmoidal) units trained using the delta rule.

The general two-layer network architecture is depicted in Figure 1 B. All networks consisted of 180 input units and 9 output units. The input units were divided into three groups. Each group of sixty units encoded the circumferential profile for each view. A relative coding scheme was used in which the sixty ordered segments constituting each view were divided by the segment with the highest scintigraphic count. Each input pattern was then normalized to unit vector

[1] Also sometimes referred to as LAO 70 or the "steep LAO" view.

[2] Pre-processing involved positioning of the region of interest (ROI), interpolative background subtraction, smoothing and rotational alignment to the heart's apex [8].

[3] The profiles were generated using the Elscint CTL software package for planar quantitative thallium-201 based on the Cedars-Sinai technique [9, 13, 1].

length. Target numerical values were assigned to the categorical visual scores to make the data suitable for network learning: *normal* = 0.0, *mild defect* = 0.3, *moderate defect* = 0.7, and *severe defect* = 1.0.

2.1.1 Weight smoothing

Weight smoothing "penalized" adjacent fan-in weights to the hidden units in proportion to their difference:

$$C_{ij} = (w_{ij} - w_{i,j-1}) + (w_{ij} - w_{i,j+1}) \tag{1}$$

where $j-1$ and $j+1$ are the left and the right "neighbors" of j. This rule is a special case of one proposed elsewhere [11]. The value of C_{ij} was subtracted from Δw_{ij} on each weight update:

$$w_{ij} = w_{ij} + \Delta w_{ij} - \beta C_{ij} \tag{2}$$

In the present experiments, the smoothing parameter β, was set to 0.01.

2.1.2 Radial Basis Function Networks

RBF networks [14, 15, 3] typically consist of two layers of units: one layer of radial units, typically gaussian, and a second layer of semi-linear units. In our experiments, the activation value of a gaussian unit, O_j, is given by:

$$net_j = \sum_i (w_{ij} - \frac{v_i}{||v||})^2 \tag{3}$$

$$O_j = exp(-\frac{net_j}{\omega}) \tag{4}$$

where j is an index to a gaussian unit, i is an input unit index, and $||v||$ is the length of the current input vector v (for a single view). The width of the gaussian, given by ω, was fixed at 0.25 for all units.

The gaussian units were trained using a competitive learning scheme that moves the center of the unit closest to the current input pattern (O_{max}, i.e. the "winner") closer to the input pattern:

$$\Delta w_{i,winner} = \eta(\frac{v_i}{||v||} - w_{i,winner}) \tag{5}$$

In addition, the other units were also pulled towards the input pattern, although to a much smaller extent[4]. We used a ratio of 1 in 100.

The second layer was trained using the delta-rule [20]. Input to the second layer consisted of the activations of the entire set of nine gaussian units plus an additional unit per view to encode the scaling factor of the input patterns lost as a result of input vector normalization. This additional unit was clamped to the minimum value of the input pattern for that view[5].

In our experiments, the competitive and the supervised layers were trained sequentially: first competitive then supervised.

2.1.3 Training

Because of the limited number of cases available, we employed the "leave-one-out" or "jack-knife" method of training. We trained 100 separate networks; each network was trained on a subset of 99 of the 100 cases and tested on the remaining one. This procedure enabled us to determine the generalization performance for each case.

Each of the 100 networks was initialized in the normal way; Weights were initialized to small random values between 0.5 and -0.5 and trained with learning rate (η) = 0.05 and momentum (α) = 0.9.

[4] Referred to as the *leaky learning model* in [18]

[5] The scaling factor was given as additional input to the output units in the other networks as well to ensure that the that the comparisons among them were fair.