CLINICAL DIAGNOSIS BASED ON DEEP LEARNING

PEDRO SAAVEDRA

1. Introduction

> *Recognizing that machines in the practice of medicine are here to stay, physicians have the obligation to learn as much of their advantages and limitations as they can comprehend* (1).

This phrase appears in an article by Keeve Brodman, published in 1959. Since then, the so-called ***machine learning*** methods have become increasingly present in the analysis of biomedical data, particularly in the field of clinical diagnostics This is part of the so-called ***supervised learning***, the purpose of which is to construct classifiers based on a set of predictors, including images of the patients. In this sense, artificial neural networks (ANN) set a precedent since the 1940s (2), but became very promising in the late 1980s (3). Subsequently, support vector machines, boosting, classification and regression trees emerged, and ANN fell into disuse, partly because they required costly adjustments, while the new methods were more automatic. Moreover, on many problems the new methods outperformed poorly trained neural networks. This was the status quo during the first decade of the new millennium. The use of neural networks resurfaced starting in 2010 with the new name of ***deep learning***, new architectures, more features, and a number of successes on some specific problems, such as image and video classification and speech and text modeling. Many experts believe that the main reason for these successes is the availability of ever larger training datasets, made possible by the large-scale use of digitization in science and industry and, no doubt, ever better computing resources.

The purpose of this paper is to make an approach to ANNs based on logistic regression models, which can be considered as very elementary ANN models. We illustrate these models using the well-known Candida-Score classifier (4, 5) designed to predict invasive candidiasis in critically ill patients with fungal colonization. In third section we introduce the basic concepts of the ANNs. After a brief description of the array structure

of the images, we introduce the convolutional neural networks (CNNs). Finally, we trained a CNN for the diagnosis of pneumonias based on Chest X-ray.

2. Prediction of invasive candidiasis by candida-score

We will make a heuristic approach to the neural network concept by illustrating it with the study of the *candida score* (5) summarized below.

Objective: To obtain a score for deciding early antifungal treatment when candidal infection is suspected in nonneutropenic critically ill patients.

Design: Analysis of data collected from the database of the EPCAN project, an ongoing prospective, cohort, observational, multicenter surveillance study of fungal infection and colonization in intensive care unit (ICU) patients.

Interventions: Surveillance cultures of urine, tracheal, and gastric samples were obtained weekly. Of the 980 patients with candida colonization, 97 had proven candidal infection.

Statistical analysis: The odds ratio (OR) for each risk factor associated with colonization vs. proven candidal infection was estimated. A logistic regression model was performed to adjust for possible confounders. The "Candida score" was obtained according to the logit method.

In Table 1, categorical variables are expressed as frequencies and percentages and continuous as mean and standard deviation (SD) when data followed a normal distribution, or as median and interquartile range (IQR = 25th - 75th percentile) when distribution departed from normality. The percentages were compared, as appropriate, using the Chi-square (χ^2) test or the exact Fisher test, the means by the t-test, and the medians by the Wilcoxon test for independent data.

Candida Score is a diagnosis rule defined by means of a logistic model of the form:

$$\Pr\left(IC = 1 \mid Sepsis, TPN, MF, Surgery\right) =$$

$$\frac{\exp\left(\beta_0 + \beta_1 Sepsis + \beta_2 TPN + \beta_3 MF + \beta_4 Surgery\right)}{1 + \exp\left(\beta_0 + \beta_1 Sepsis + \beta_2 TPN + \beta_3 MF + \beta_4 Surgery\right)}$$

TABLE 1. Description of variables in EPCAN project.

	Overall: N = 980	No IC : N = 883	IC : N = 97	P-value
Age	58.9 ± 17.0	58.9 ± 17.0	58.5 ± 16.9	0.825
Apache	19.0 ± 7.7	19.1 ± 7.8	18.2 ± 6.5	0.273
BMI	26.3 ± 4.9	26.3 ± 4.6	26.3 ± 6.6	0.99
Sex.male	642 (65.7)	574 (65.2)	68 (70.1)	0.337
DM	150 (15.3)	136 (15.4)	14 (14.4)	0.801
COPD	211 (21.5)	197 (22.3)	14 (14.4)	0.073
Neoplasia	103 (10.5)	85 (9.6)	18 (18.6)	0.006
Hematopathy	20 (2.0)	17 (1.9)	3 (3.1)	0.439
HIV	16 (1.6)	14 (1.6)	2 (2.1)	0.667
Cirrhosis	42 (4.3)	40 (4.5)	2 (2.1)	0.424
Renal.failure	48 (4.9)	44 (5.0)	4 (4.1)	1
Mechanical.Ventilation	929 (94.8)	837 (94.8)	92 (94.8)	0.982
Enteral.Nutrition	763 (77.9)	695 (78.7)	68 (70.1)	0.053
SEPSIS	219 (22.4)	156 (17.7)	63 (65.0)	< 0.001
Parenteral.Nutrition	547 (55.8)	462 (52.3)	85 (87.6)	< 0.001
MF	563 (57.5)	494 (56.0)	69 (71.1)	0.004
SURGERY	309 (31.5)	258 (29.2)	51 (52.6)	< 0.001
SDD	60 (6.1)	50 (5.7)	10 (10.3)	0.07
Death	411 (42.5)	355 (40.8)	56 (57.7)	0.001
Patient.Type				< 0.001
Medical	483 (49.3)	449 (50.8)	34 (35.1)	
Surgical	309 (31.5)	258 (29.2)	51 (52.6)	
Traumatology	188 (19.2)	176 (19.9)	12 (12.4)	
ICU.days	21 (14; 33)	20 (14; 33)	28 (17; 42)	< 0.001
Hospital.days	39 (25; 63)	38 (25; 63)	46 (27; 61)	0.176

[a] Data are means (SD), frequencies (%) and medians (Percentiles 25-75)

To estimate this model, the data set was spplited into a *learning data* set or *training data* ($n = 627$) and a *test data* set ($n = 353$). The estimation of the logistic model based on the learning data is shown in Table 2.

Thus, Candida Score is defined as:

$$CS = \beta_1 Sepsis + \beta_2 TPN + \beta_3 MF + \beta_4 Surgery$$

Note that:

TABLE 2. Estimation of the logistic model for the invasive candidiasis (Candida-Score).

	Coefficient (SE)	P-value	AIC *	Odd-Ratio (95% CI)
(Intercept)	-4.916 (0.485)			
SEPSIS	2.038 (0.314)	< 0.001	349.6	7.678 (4.145 ; 14.221)
TPN	0.908 (0.389)	0.014	310.7	2.478 (1.157 ; 5.309)
MF	1.113 (0.379)	0.002	314.6	3.044 (1.448 ; 6.397)
SURGERY	0.997 (0.319)	0.002	314.6	2.709 (1.450 ; 5.063)

[a] (*) AIC if the factor is removed; AIC is a measure of lack of fit: For the full model, AIC = 300.6. Note that if any of the factors are removed, the AIC value is increased, thus worsening the fit.

$$\Pr\left(IC = 1 \mid Sepsis, TPN, MF, Surgery\right) = \frac{\exp\left(\beta_0 + CS\right)}{1 + \exp\left(\beta_0 + CS\right)}$$

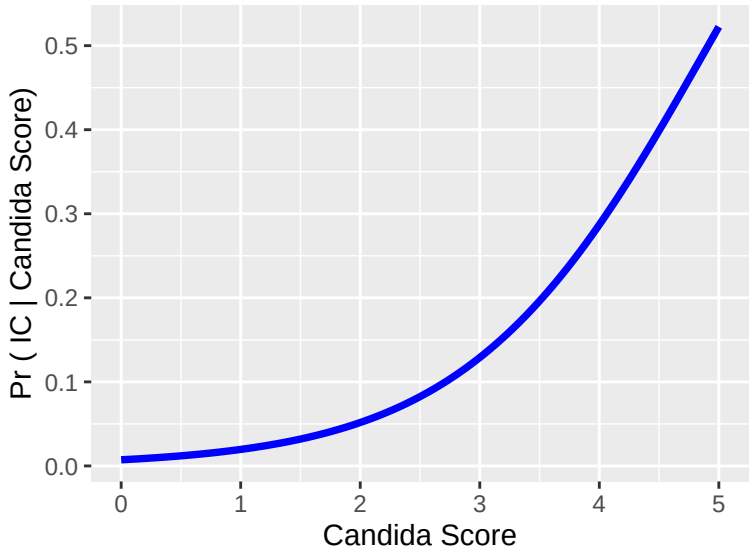

FIGURE 1. Probability of invasive candidiasis based on Candida score value.

Figure 2 shows a schematic representation of the candida score. Such a graph suggests the idea of biological neural networks.

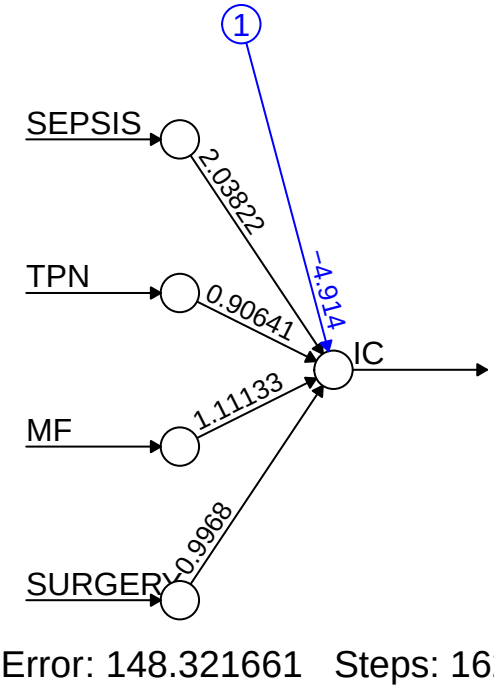

FIGURE 2. Candida-Score represented in the form of artificial neural network .

3. Deep learning: introduction to the neural networks

Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy. Deep learning is a subset of machine learning, which is basically a neural network with three or more layers. These neural networks attempt to emulate the behaviour of the human brain -though far from matching its capabilities- but allow it to "learn" from large amounts of data. Although a neural network with a single layer can already make approximate predictions, additional hidden layers help to optimize and refine accuracy.

The purpose of an *artificial neural network* (ANN) is to obtain a predictive rule (classifier) based on a feature variables observed on each patient.Formally, we consider a given classification $\mathcal{C}$ (e.g., alternative diagnoses corresponding to the disease), and several feature variables $X_1, ..., X_p$ (e.g., a set of clinical variables with information about the study disease) whose purpose is to predict the most likely class. The predictor has the form:

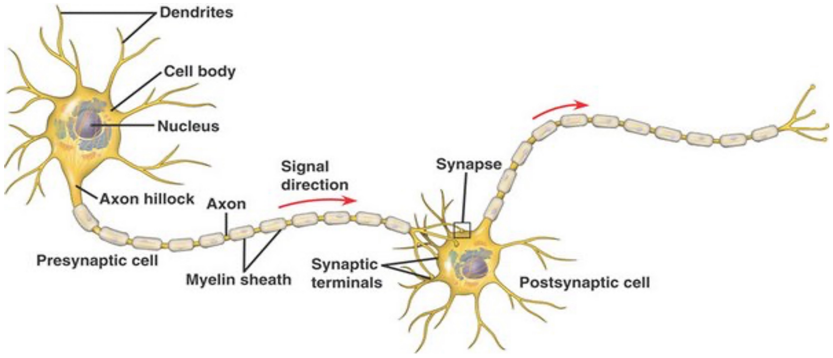

FIGURE 3. A biological neuron is a special cell consisting of a nucleus, a body, and connectors. Each neuron has a close connection with thousands of other neurons. Electrochemical impulses are transmitted through this connection, causing the entire neural network to be in a state of excitation or viceversa.

$$\Pr\left(diagnosis = d \mid X_1, ..., X_p\right) : d \in \mathcal{C}$$

Neural network: input with the markers of the CS and a hidden layer with two neurons. In order to improve the prediction of invasive candidiasis, an attractive idea would be to interpolate between the input and output layers of the candida score (Figure 2) a hidden layer, for example with two neurons.

The artificial neural network shown in Figure 4 has the same characteristic variables as input, but with a hidden layer with two neurons. Note that the model from which the candida score is derived is based on five parameters while the model with two hidden layers raises the number of parameters to 12. In addition, the simple interpretability of the CS vanishes.

The elements of the neural network are:

A. Input.

$$\{SEPSIS, TPN, MF, SURGERY\}$$

B. Hidden layer.** The hidden layer is defined by two neurons, each a function of the input variables; namely

$$Z_1 = \sigma\left(1.94 + 6.37 \cdot SEPSIS + 1.13 \cdot TPN - 7.35 \cdot MF - 0.84 \cdot SURGERY\right)$$

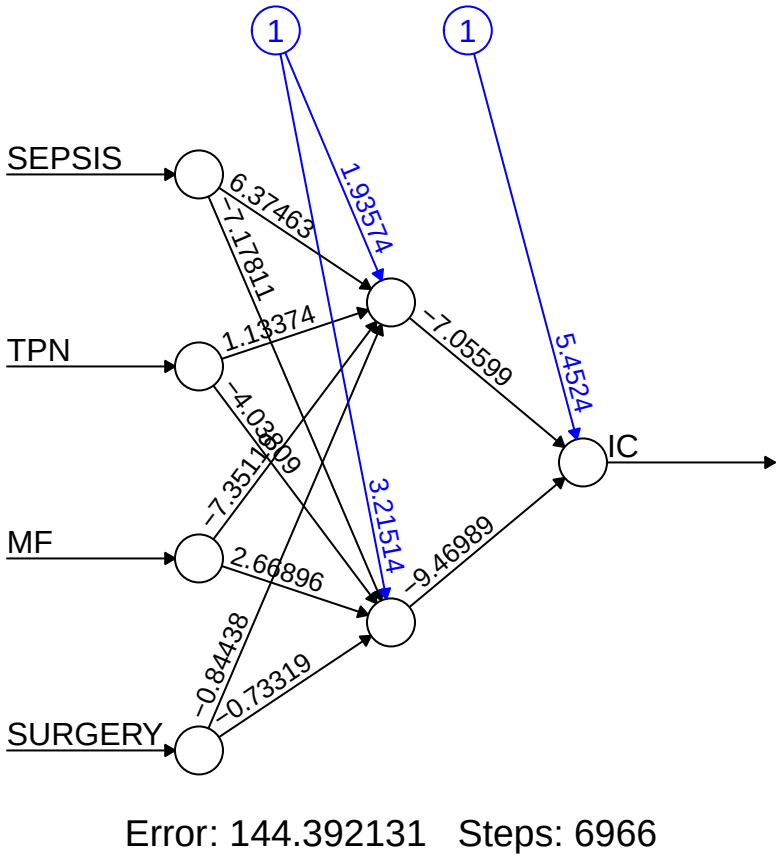

FIGURE 4. Artificial neural network based on the same variables as candida-score but with a hidden layer of two neurons.

and

$$Z_2 = \sigma\left(3.22 - 7.18 \cdot SEPSIS - 4.04 \cdot TPN + 2.67 \cdot MF - 0.73 \cdot SURGERY\right)$$

Here, $\sigma(.)$ is the called *activation function*, which generally is (*sigmoid function*):

$$\sigma\left(v\right) = \frac{1}{1 + e^{-v}}$$

C. Output layer: prediction of the *invasive candidiasis* from input variables. From the neurons Z_1 and Z_2, we obtain the variable T as:

$$T = 5.45 - 7.06 \cdot Z_1 - 9.47 \cdot Z_2$$

Finally, the following transformation, in addition to making the model more flexible, allows us to obtain a predictor with values between 0 and 1.

$$\Pr\left(IC = 1 \mid SEPSIS, TPN, MF, SURGERY\right) = \frac{\exp\left(T\right)}{1 + \exp\left(T\right)}$$

We can try to improve the prediction model by adding variables that showed statistical significance in the univariate analysis. To this end, we now consider the data set:

$$\{(SEPSIS_i, TPN_i, MF_i, SURGERY_i\, Neoplasia_i, ICU.days, ; IC_i) \ : \ i = 1, ..., 980\}$$

We also trained the corresponding neural network using the same 627 training data. Figure 6 shows that ANN based on these six variables does not perform a better classification task than logistic regression.

4. Medical diagnosis based on images

Structures of the images. Basically, a digitized image is a three-dimensional matrix. The width and height correspond to pixels, which are encoded on an integer scale from 0 to 255. These values are rescaled to the interval [0,1] by dividing by 256. The depth (canal) is three units and corresponds to the three basic colors red, green and blue. An image in shades of gray is reduced to the two-dimensional pixel array.

Supervised classification based on images. For a given classification $\mathcal{C}$, the purpose of a CNN is to obtain a predictive rule (classifier) based on a patient image, which we denote by $\mathcal{I}mage$. The classifier is a function that assigns to each image a probability distribution over $\mathcal{C}$ of the form:

$$\Pr\left(diagnosis = d \mid \mathcal{I}mage\right) : d \in \mathcal{C}$$

The patient is then assigned the most likely diagnosis (class). Note that:

$$\sum_{d \in \mathcal{C}} \Pr\left(diagnosis = d \mid \mathcal{I}mage\right) = 1$$

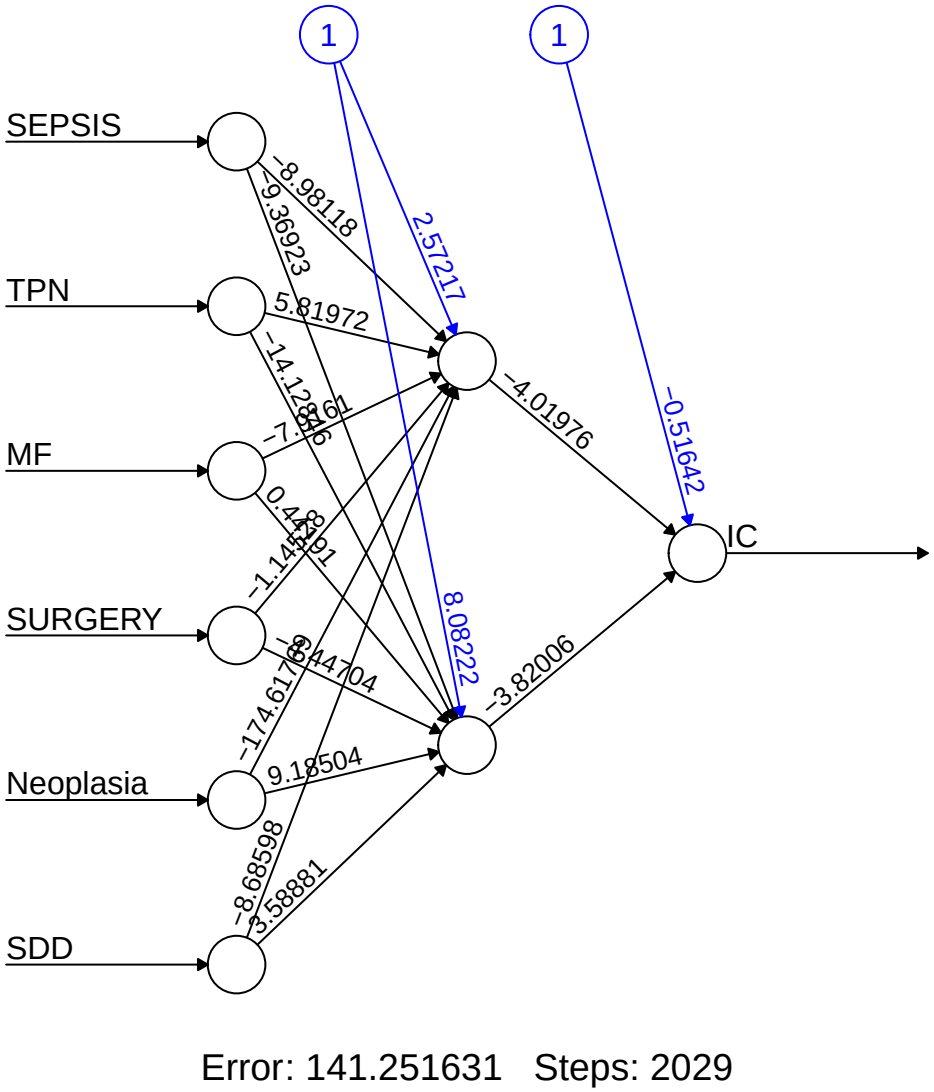

FIGURE 5. ANN to predict invasive candidiasis based on those variables that in univariate analysis showed a significant association with the outcome at $P < 0.1$.

5. Convolutional neural networks (CNN)

A special family of neural networks are the convolutional neural networks (CNNs) which are mainly used for classifying images and have shown spectacular success on a wide range of problems. CNNs mimic to some degree how humans classify images, by recognizing specific features or patterns anywhere in the image. The network first identifies low-level features in the input image, such as small edges, patches of color, and the like. These low-level features are then combined to form higher-level features, such as

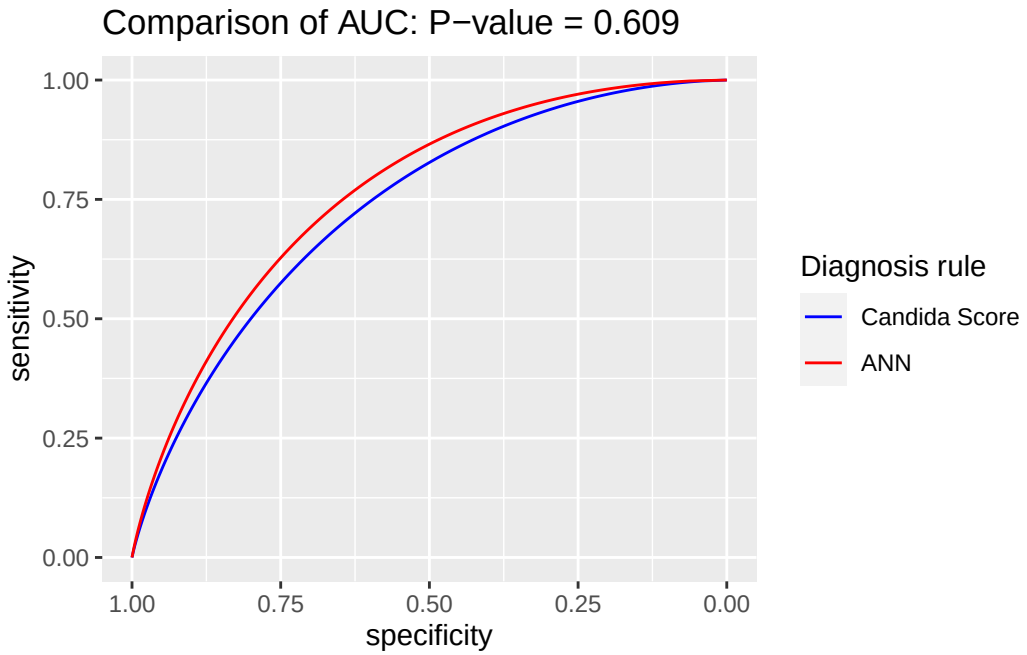

FIGURE 6. Predictors of invasive candidiasis. Comparison of ROC curves for CS and ANN based on six variables.

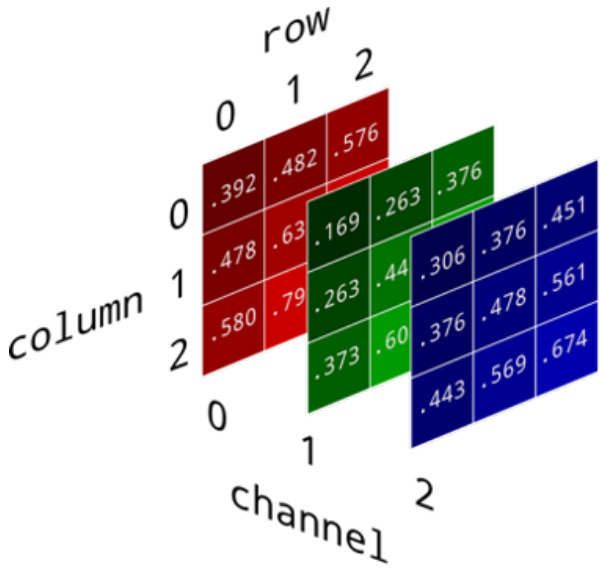

FIGURE 7. A colour image is a three-dimensional array. The first two dimensions correspond to the pixel map and the third to the channel. When the image is in grey tones, it is reduced to a two-dimensional array (a single channel).

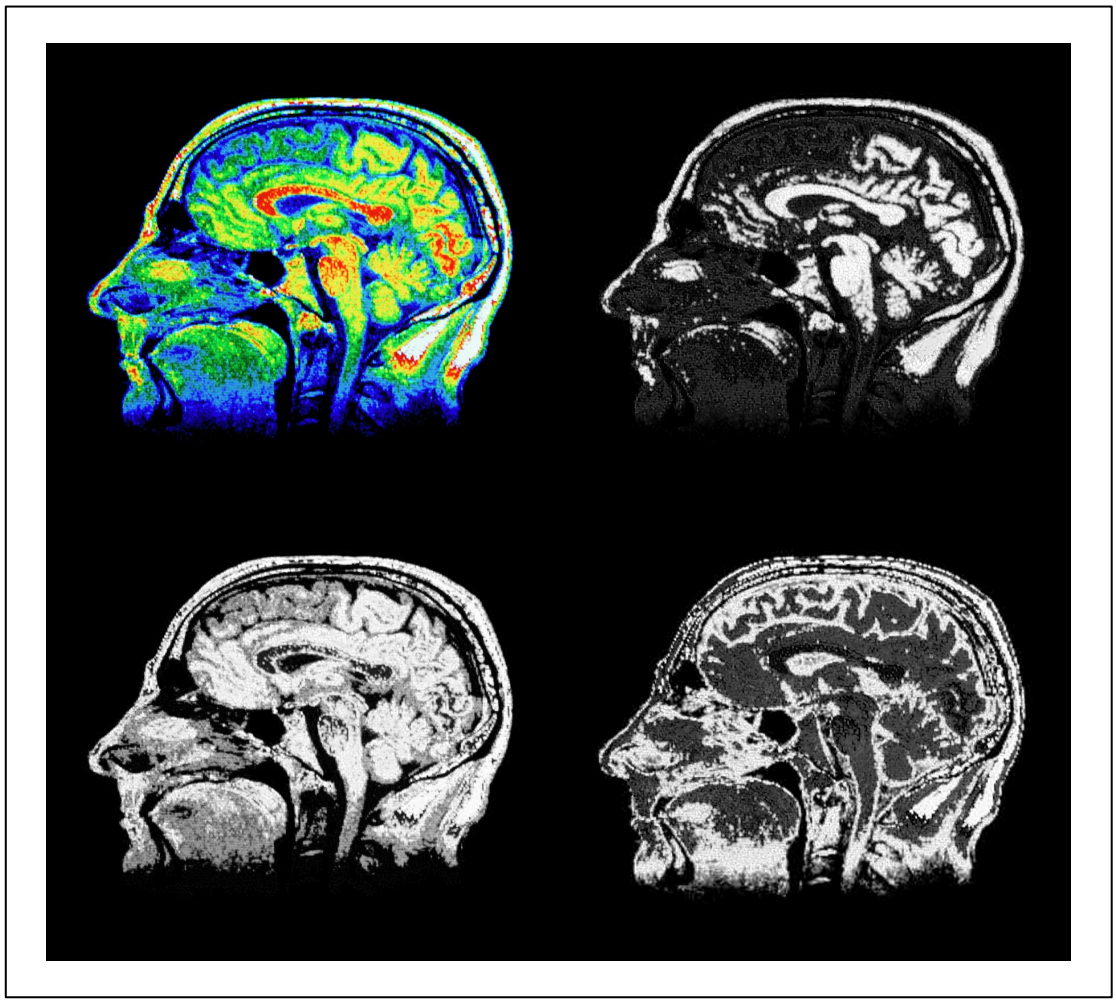

FIGURE 8. Magnetic resonance imaging

parts of ears, eyes, and so on. Eventually, the presence or absence of these higher-level features contributes to the probability of any given output class.

This process is known as the image preprocessing, which has essentially two phases, namely ***filtering*** and ***pooling***. We now describe each one of these phases.

Convolutional layer. A convolution layer is made up of a large number of convolution filters, each one being a template that determines whether a particular local feature is present in an image. To fix ideas, consider the following image 4×3:

$$Input\ image = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \\ j & k & l \end{bmatrix}$$

and the 2×2 filter.

$$Filter = \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$

The convoluted image obtained is then:

$$Convolved\ image = \begin{bmatrix} a\alpha + b\beta + d\gamma + e\delta & b\alpha + c\beta + e\gamma + f\delta \\ d\alpha + e\beta + g\gamma + h\delta & e\alpha + f\beta + h\gamma + i\delta \\ g\alpha + h\beta + j\gamma + k\delta & h\alpha + i\beta + k\gamma + l\delta \end{bmatrix}$$

To highlight the edges of the glioma image (fig), we can apply the following filter:

$$HB = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Pooling layer. A pooling layer is a procedure whose purpose is to condense a large image into a smaller pooling summary image. While there are a number of possible ways to perform pooling, the max pooling operation summarizes each non-overlapping 2×2 block of pixels in an image using the maximum value in the block. This reduces the size of the image by a factor of two in each direction. Here is a simple example:

$$\max.pool \begin{bmatrix} 4 & 2 & 1 & 4 \\ 9 & 5 & 5 & 4 \\ 8 & 0 & 3 & 2 \\ 2 & 2 & 3 & 5 \end{bmatrix} = \begin{bmatrix} 9 & 5 \\ 8 & 5 \end{bmatrix}$$

Architecture of a Convolutional Neural Network. For an image-based classification (diagnosis), the architecture of a convolutional neural network (CNN) basically consists of the following elements:

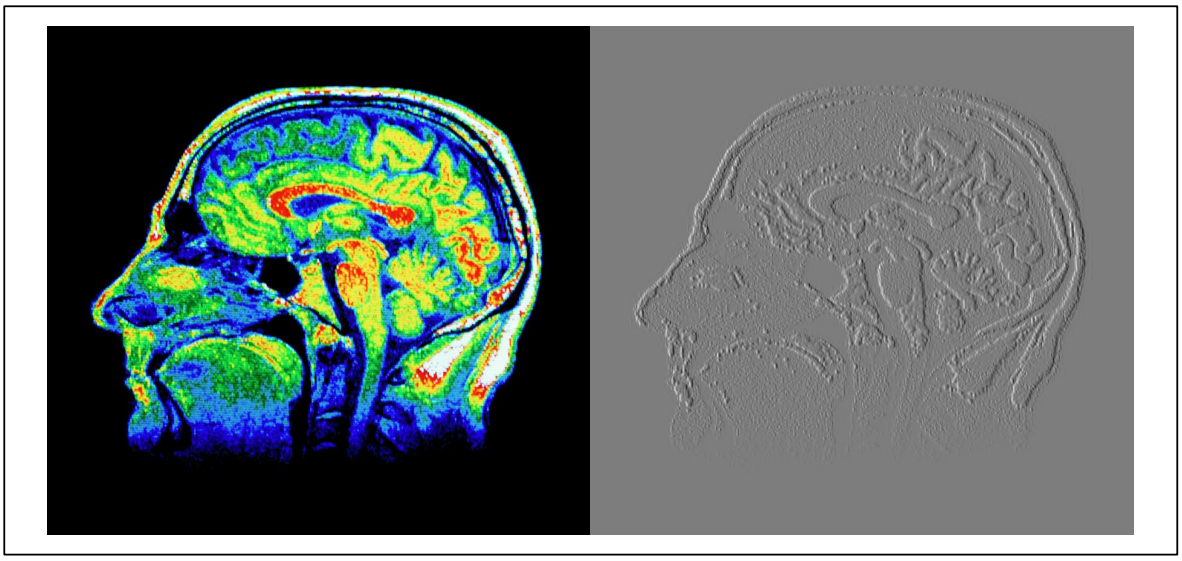

Magnetic resonance imaging: original (900 x 900 x 3) and Sobel filter

FIGURE 9. Original image and Sobel filter

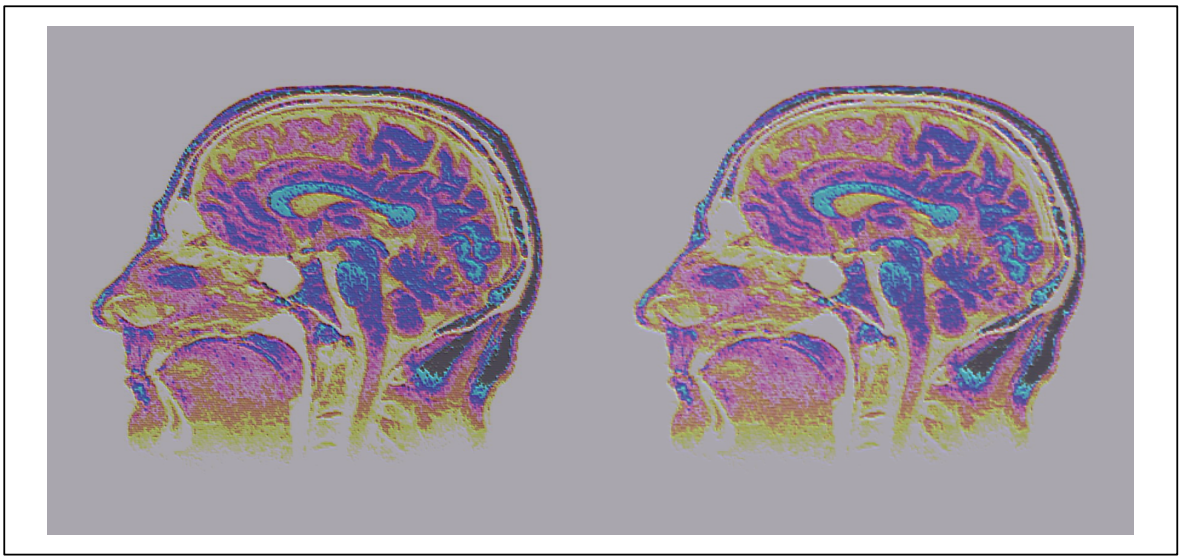

Magnetic resonance imaging: filtered (898 x 898 x 3) and Max−Pooling (499 x 499 x 3)

FIGURE 10. Original image and max-pooling

1. The input, which is an image determined by an array of either $M \times N$ pixels (gray shades only) or three arrays of the same dimension corresponding to the colors red, green and blue (dimension $M \times N \times 3$). The input image is labeled with each of the possible diagnoses (number of categories of the classification variable).

2. A sequence of pairs of layers corresponding each pair to K convolution matrices, which extract features relevant for classification and the subsequent K matrices obtained from max-pooling, which reduce the sizes of the previous matrices while keeping the relevant information. The latter can be considered as a single multi-channel image (K channels), which is the input for the next layer.

3. A subsequent fully connected ANN whose input is the result of the flattening of successive pixel arrays.

4. Finally, a probability distribution over the classes considered (alternative diagnoses) given by the ANN.

Each filtering (convolution) of the features of a multi-channel image (K channels) is obtained by applying a K-channel filter, thus obtaining a multi-channel image. The K channels are then summed to form a two-dimensional feature map. The elements of the filter-matrices are unknown parameters that are estimated in the CNN training process. It can therefore be said that the filters are learned from the training data.

6. Classification of pneumonias in a pediatric population based on chest X-rays

We illustrate the training process of a convolutional neural network using a set of 5,840 300x300 pixel chest X-rays available to generate a classifier to distinguish between healthy lung, bacterial pneumonia and viral pneumonia.In the following points we summarise the objective of the problem and the construction of the classifier.

Purpose: Discriminate between non-pneumonia, bacterial pneumonia and viral pneumonia using chest X-rays.

Data. A total of 5,840 (300×300 pxs) images are used, distributed as follows:

Method. Training of a convolutional neural network with the 5,216 images (data training) and validation of the classifier with 624 data test.

TABLE 3. Original image and max-pooling

	No	Bacterial	Viral
Data training	1341	2530	1345
Data test	234	242	148

[a] Distribution of the 5,840 images of dimension 300x300

Model. The classifier gives the probability of each diagnosis as:

$$\Pr\left(Diagnosis = d \mid chestX - ray\right) \ : \ d \in \{Normal,\ Bact\,, Viral\}$$

Validation. Each subject is assigned to the most probable class. Goodness of fit was evaluated as the rate of well-classified observations using the test data.

Data. The distribution of the 5,440 chest radiographs according to gold-standard diagnosis and type of data is shown in the following table

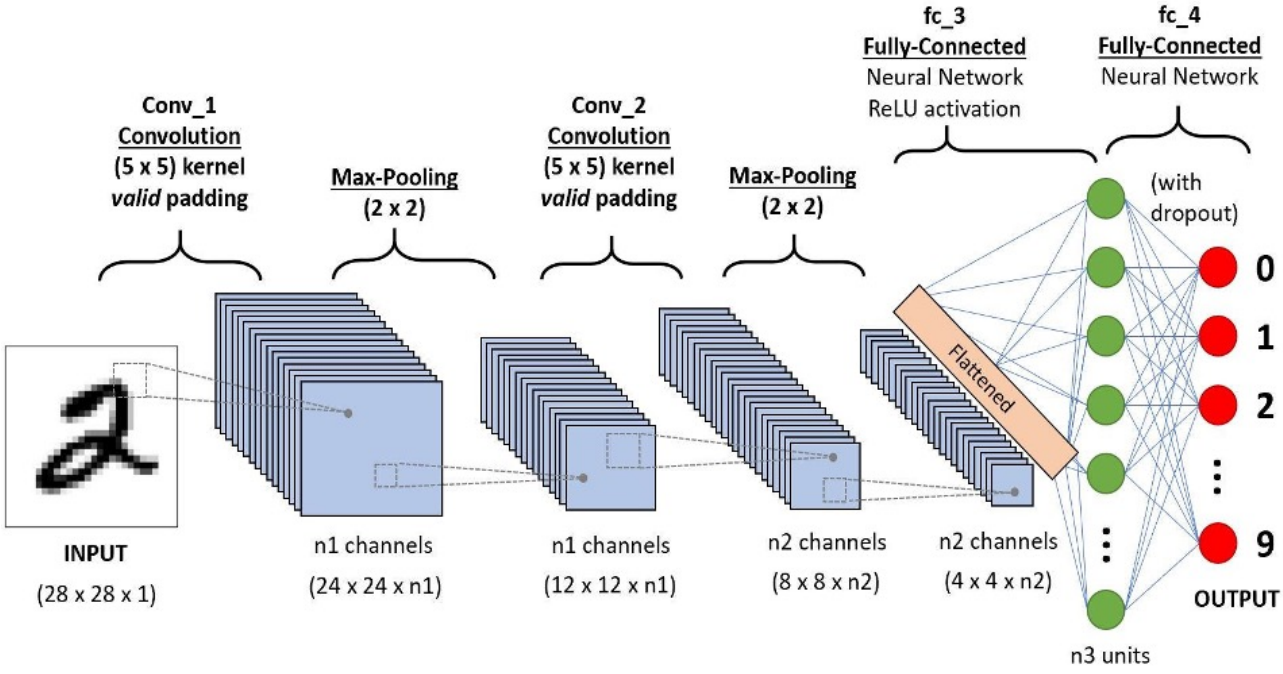

FIGURE 11. Convolutional Neural Network Architecture.

Table 4 shows the architecture of the convolutional network for the diagnosis of pneumonia.

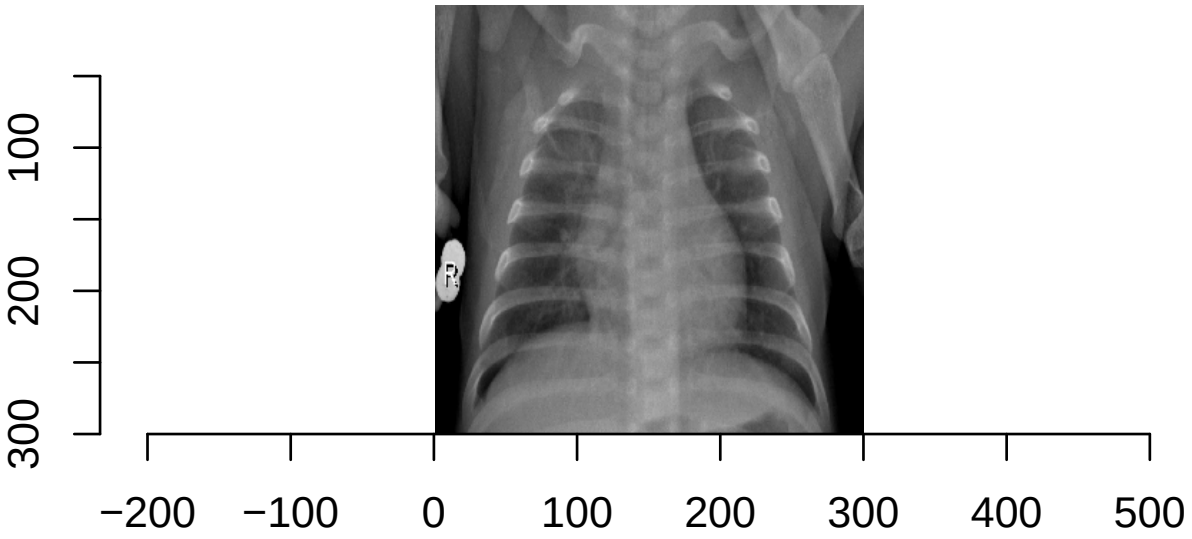

FIGURE 12. chest X-rays 300 x 300 pxs

TABLE 4. Convolutional neural network architecture for the diagnosis of pneumonias.

Layer (type)	Output Shape	Param #
conv2d_2 (Conv2D)	(None, 298, 298, 32)	320
max_pooling2d_2 (MaxPooling2D)	(None, 149, 149, 32)	0
conv2d_1 (Conv2D)	(None, 147, 147, 64)	18496
max_pooling2d_1 (MaxPooling2D)	(None, 73, 73, 64)	0
conv2d (Conv2D)	(None, 71, 71, 128)	73856
max_pooling2d (MaxPooling2D)	(None, 35, 35, 128)	0
dropout_1 (Dropout)	(None, 35, 35, 128)	0
flatten (Flatten)	(None, 156800)	0
dense_1 (Dense)	(None, 128)	20070528
dropout (Dropout)	(None, 128)	0
dense (Dense)	(None, 3)	387
		Total params: 20,163,587 Trainable params: 20,163,587 Non-trainable params: 0

TABLE 5. Gold standard and prediction by CNN.

PNeumonia	Diagnosis	No	Bacterial	Viral
Gold.Standard	No pneum.	105	75	54
	Bacterial pneum.	2	237	3
	Viral pneum.	0	51	97

Prediction using datatest. Table 5 compares the gold standard of diagnosis with the prediction obtained using the convolutional neural network. Note rate of well-classified observations was 70.4%. and the low false negative rate.

7. References

[1] Brodman K, Van Woerkom AJ, Erdmann AJ Jr, Goldstein LS. Interpretation of symptoms with a data-processing machine. A.M.A. Arch Intern Med 1959;103:776–82. https://pubmed.ncbi.nlm.nih.gov/13636498.

[2] McCulloch WS, Pitts WA logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysics 5, 115–133 (1943). https://doi.org/10.1007/BF02478259.

[3] James, G., Witten,D., Hastie,T. and Tibshirani,R. An Introduction to Statistical Learning with Applications in R. New York, Springer Science and Business Media, 2013, 41.98, eISBN:978-1-4614-7137-7. https://doi.org/10.1007/978-1-0716-1418-1

[4] Chollet, F., and J. J. Allaire. "Deep Learning with R Manning Publications Co." Shelter Island, NY (2018).

[5] León C, Ruiz-Santana S, Saavedra P, Almirante B, Nolla-Salas J, Alvarez-Lerma F, Garnacho-Montero J, León MA; EPCAN Study Group. A bedside scoring system ("Candida score") for early antifungal treatment in nonneutropenic critically ill patients with Candida colonization. Crit Care Med. 2006 Mar;34(3):730-7. https://pubmed.ncbi.nlm.nih.gov/16505659.