

# Un modelo nuevo variacional para la clasificación binaria en el contexto de aprendizaje automático supervisado

C. D. Brito Pacheco\*, C. F. Brito Loeza

Facultad de Matemáticas, Universidad Autónoma de Yucatán, Anillo Periférico Norte, Tablaje Cat., 13615, Mérida Yucatán México.

Fecha de recepción: 16 de mayo de 2018 — Fecha de aceptación: 21 de agosto de 2018

## Resumen

Se examina el problema de aprendizaje supervisado en su planteación continua. Posteriormente se da una condición de optimalidad general a través de técnicas del análisis funcional y el cálculo de variaciones. Esto nos permite resolver la condición de optimalidad para la función deseada  $u$  numéricamente y hacer varias comparaciones con otros modelos de aprendizaje supervisado ampliamente utilizados. Se emplea la precisión porcentual y el área bajo la curva característica operativa del receptor (AUC por sus siglas en inglés) como métricas del rendimiento. Finalmente, se realizan 3 análisis basados en estas dos métricas mencionadas donde comparamos los modelos y hacemos conclusiones para determinar si nuestro método es competitivo o no.

**Palabras Clave:** aprendizaje variacional, aprendizaje automático, inteligencia artificial

## A New Variational Model for Binary Classification in the Supervised Learning Context

### Abstract

We examine the supervised learning problem in its continuous setting and give a general optimality condition through techniques of functional analysis and the calculus of variations. This enables us to solve the optimality condition for the desired function  $u$  numerically and make several comparisons with other widely utilized supervised learning models. We employ the accuracy and area under the receiver operating characteristic curve as metrics of the performance. Finally 3 analyses are conducted based on these two mentioned metrics where we compare the models and make conclusions to determine whether or not our method is competitive.

**Keywords:** variational learning, supervised learning, artificial intelligence.

### Introduction

The problem of *supervised learning* arises in contexts where a study of a dependent variable  $Y$  is necessary in terms of an independent variable  $X$ . The goal of supervised learning is to predict the values of  $Y$  given many instances of the variable  $X$ . Generally, we call the instances of  $X$  and the values  $Y$  takes, the *inputs* and *outputs*. In the pattern recognition context we typically may read *features* and *responses* respectively as an alternative terminology. With this in mind, it is noteworthy that we are in fact assuming the existence of a function between

$X$  and  $Y$  such that  $Y = f(X)$ . That is, there exists a relationship between the two variables. In general, we say there are two main tasks in supervised learning: *classification* and *regression*. We make the distinction between these two contexts by defining classification as the *task that classifies inputs into several classes* and regression as the *task which estimates the functional relationship between two variables*. To summarize the supervised learning concerns itself with the estimation of a function  $\hat{f}$  such that  $Y \approx \hat{f}(X)$ . Admittedly, the interpretation of supervised learning is to let a mathematical or computational model “learn by example”. The examples being

\*carlosdotb@gmail.com

the values  $X$  takes, while the learning process is the identification of patterns given by the relationship between  $X$  and  $Y$ . Chiefly, supervised learning is catalogued as a form of *artificial intelligence* [1]. In recent years many supervised learning methods have been developed and many have been reviewed time and time again [2–5]. One of the most recent reviews was conducted by Singh *et al.* [2], which tests many different methods on a single dataset and presents the advantages and disadvantages in a textual manner and establishes their application areas. A slightly older but extremely exhaustive analysis was conducted by Fern'andez-Delgado *et al.* which evaluates 179 classifiers arising from 17 families on 121 datasets [3]. Support Vector Machines (SVB) and Random Forests are the top performing methods. Kotsiantis *et al.* [4] have conducted a more comprehensive review of various different classification methods in which they show that SVM and Neural Networks (NN) were the best performing algorithms in terms of accuracy, classification speed and tolerance to parity problems. These algorithms, however, are lacking in terms of speed of learning, danger of overfitting and model parameter handling where Naïve Bayes and kNN are the top contenders. After analysing all of the results more carefully, one may argue that SVM clearly outperforms NN. This being said, one will find that NN and SVM are the most used supervised learning methods due to their simple “black box” functionality and so make good candidates to compare against. This work will not be concerned with such analyses but will instead propose a new supervised binary classification method. The method will be based in the resolution of a mathematical model through techniques from calculus of variations and functional analysis. Employment of these techniques allows us to exploit the underlying mathematical theory and properties from each of these fields. Even more, by utilizing these techniques we may obtain meaningful conditions in the continuous sense rather than performing optimization discretely. Comparing the performance of the model with SVM and NN we will determine if the new method is competitive or not.

The structure of this work is as follows, the section “A Variational Approach for Supervised Learning” will consist in giving the formal problem statement. We will discuss some of the characteristics of this problem and will proceed to give a solution. Furthermore, we will present choices of loss and regularization functions to generate a new model from our solution. Next, we will present “The Completed Model” in the section by the same name and discuss its training method. Moving on, we evaluate the model and conduct analyses on the results in the sections “Evaluation Methodology” and “Results and Analysis”. Lastly, we shall finalize this work in the section “Conclusion”.

## A Variational Approach for Supervised Learning

### The Problem

The supervised learning problem may be characterized as follows. Given a training set of  $N$  observations

$$\mathcal{T} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N) \mid \mathbf{x}_i \in \Omega, y_i \in \mathcal{Y}\},$$

we would like to find a function  $u = u(\mathbf{x})$  defined by  $u : \Omega \rightarrow \mathcal{Y}$ . In order to find a suitable function  $u$  it is necessary to suggest a suitable model which best describes the problem. The most widely used framework to solve the supervised learning problem is the minimization of the sum of a loss function  $L(u, y)$  with an added regularization term  $S(u)$ :

$$\min_u \lambda S(u) + \sum_{i=1}^N L(u(\mathbf{x}_i), y_i).$$

For the continuous case, however, we have:

$$\min_u \lambda S(u) + \int_{\Omega} L(u(\mathbf{x}_i), y_i). \quad (1)$$

The idea behind the model is to “punish” incorrect predictions using the loss function while the regularizer term allows us to “control” or “regulate” the variable  $u = u(\mathbf{x})$ . A regularization term is typically used to solve ill-posed problems or prevent overfitting. Thus, the regularization term should only depend on the variable being regularized; in our case  $u$ . We may further specialize  $S(u)$  to our needs. Opting to punish the complexity of  $u$ , we might attempt to “measure” the complexity and subsequently punishing it. Thereupon we let  $R(u) > 0$  be a Lebesgue integrable function which measures the complexity of  $u$ . The exact definition of  $R$  is left until a later section. This being said, we substitute back into Equation 1 to obtain:

$$\min_u \int_{\Omega} [L(u, y) + \lambda R(u)] d\mathbf{x}. \quad (2)$$

This last equation will be our main interest for the remainder of this work.

One way to interpret Equation 2 is to notice that the *empirical risk* of choosing  $u$  among a space of *hypothesis functions*  $\mathcal{H}$  is defined as the expectation of the loss function  $L$ :

$$\mathbf{E}[L(u, y)] = \int_{\Omega} L(u, y) d\mathbf{x}$$

Thus, if we wanted to minimize the empirical risk, it would amount to solving the following minimization problem:

$$u^* = \arg \min_{u \in \mathcal{H}} \mathbf{E}[L(u, y)] \quad (3)$$

$$= \arg \min_{u \in \mathcal{H}} \int_{\Omega} L(u, y) dx, \quad (4)$$

where  $u^*$  is the *minimizer* of the empirical risk. Therefore, it is finally clear that we will be minimizing the *regularized empirical risk for a choice of  $u$* . Although Equation 3 is usually minimized through probabilistic and statistical methods, we will employ concepts from functional analysis and calculus of variations to find a minimization condition. The condition may then be solved numerically to obtain  $u^*$ .

### The Solution

Akin to *differential calculus*, we aim to “differentiate” the expression in Equation 2 with respect to  $u$ . Unfortunately, the usual definition of a “derivative” is insufficient for our purpose. Instead, we recall from functional analysis that if we find the functional derivative and equate to zero

$$\frac{\delta F}{\delta u(\mathbf{x})} = 0$$

we obtain the *optimality condition*. Solving this condition numerically will give us the minimizer  $u^*$  of the functional  $F[u]$ . Before deriving the optimality condition, it is necessary to make the definition of our model more precise. Let  $f$  be a function of the form:

$$f(\mathbf{x}, u(\mathbf{x}), \nabla u(\mathbf{x})) = L(\mathbf{x}, u(\mathbf{x})) + \lambda R(\mathbf{x}, u(\mathbf{x}), \nabla u(\mathbf{x})),$$

with  $y$  fixed *s.t.*  $L(u) = L(u, y)$ . Furthermore, let  $F[u]$  be the functional depending on  $u$  *s.t.*:

$$F[u] = \int_{\Omega} f(\mathbf{x}, u(\mathbf{x}), \nabla u(\mathbf{x})) dx.$$

Then, the functional derivative with respect to  $u$  is given by the PDE Euler-Lagrange Equation:

$$\frac{\delta F}{\delta u(\mathbf{x})} = \frac{\partial f}{\partial u} - \nabla \cdot \frac{\partial f}{\partial \nabla u} \quad (5)$$

$$= \left[ \frac{\partial L}{\partial u} + \lambda \frac{\partial R}{\partial u} \right] - \nabla \cdot \left[ \frac{\partial L}{\partial \nabla u} + \lambda \frac{\partial R}{\partial \nabla u} \right] \quad (6)$$

$$= \frac{dL}{du} + \lambda \frac{dR}{du} - \lambda \nabla \cdot \frac{\partial R}{\partial \nabla u} \quad (7)$$

$$= \frac{dL}{du} + \lambda \left( \frac{dR}{du} - \nabla \cdot \frac{\partial R}{\partial \nabla u} \right), \quad (8)$$

where  $u = u(\mathbf{x})$  and  $f = f(\mathbf{x}, u, \nabla u)$ . The second line is obtained by substitution. The third line is obtained by remembering that  $L$  does not explicitly depend on

$\nabla u$  and thus it must be zero. The fourth line is the factorization of  $\lambda$ . After equating to zero, our original problem is then reduced to solving the following PDE Euler-Lagrange equation for particular choices of  $L$  and  $R$ :

$$\frac{dL}{du} + \lambda \left( \frac{dR}{du} - \nabla \cdot \frac{\partial R}{\partial \nabla u} \right) = 0. \quad (9)$$

### Radial Basis Function Approximation

In order to approximate the function  $u$ , it is required we specify what kind of form it takes. Radial basis function (RBF) approximation relies on the idea that  $u(\mathbf{x})$  and can be expressed as a weighted sum of radial basis functions  $\{\phi_i(\mathbf{x})\}$ , where the weights are the fixed parameter vector  $\mathbf{w}$ . In particular, the Gaussian RBF kernel is probably the most well known and the most widely used. For this reason only, it shall be our choice all along this work. Thus, bringing the previous ideas together, we write

$$u(\mathbf{x}) = \sum_{i=1}^N w_i \phi(\mathbf{x}_i),$$

where  $\{\phi_i(\mathbf{x})\}$  will be a set of Gaussian RBF kernels given by

$$\phi_i(\mathbf{x}) = e^{-c\|\mathbf{x}-\mathbf{x}_i\|^2}$$

Maintaining the notation we have been using throughout this work, we say that the  $\{\mathbf{x}_i\}$  are the observations of the input variable  $X$ . In addition,  $c$  is a positive constant that we are free to choose, while  $\|\cdot\|$  is the Euclidean norm. We will call  $c$  the *fitting degree* of our model. As a small note, notice that by using our approximation of  $u(\mathbf{x})$ , we are setting the *centers* of the RBFs as the observations. Intuitively, this will allow our model to make predictions by calculating the Euclidean distance between a point of which the class is known, and a new input which its class is unknown while assigning a “weight” to each point. This last step is akin to saying that some points will be more important than others when making the predictions.

### Cross Entropy

In order to correctly employ a classification context we must give an appropriate loss function for classification. To find an appropriate loss, we take a Bernoulli variable  $Y \in \{0, 1\}$ . Our goal is then to predict the target class  $y$  given an input  $u(\mathbf{x}, \mathbf{w})$ , where  $\mathbf{w}$  is a vector parameter which  $u$  depends on. From the previous section, we see that this is actually the case. To derive the loss function, we choose to maximize the likelihood that, given parameters  $\mathbf{w}$ , the model results in a prediction of the correct class for each input sample with the likelihood being a function of the parameters. In fact it is the minimization of the *cross-entropy* [6] (also called *negative*

*log-likelihood*) that effectively maximizes the likelihood. This function is given by:

$$L(\sigma(u), y) = -y \ln \sigma(u) - (1 - y) \ln (1 - \sigma(u)).$$

To use this function in our framework it is necessary to calculate its derivative with respect to  $u$ . Namely

$$\frac{dL}{du} = \sigma(u) - y .$$

**Notes.** This function is convex and so coupled with a convex regularization term our model will admit a unique solution  $u^*$ . Our prediction rule is  $g(\sigma(u))$  where  $g \in \{0, 1\}$  as opposed to  $\sigma \in (0, 1)$ . We outline that  $\sigma$  gives us the probability of classifying an input  $u$  as  $y = 1$ . To effectively “binarize” the output one must define  $g$ , for example, as  $g = 1$  on  $\sigma(u) \geq 0.5$  and  $g = 0$  on  $\sigma(u) < 0.5$ .

### Laplacian Regularization $-\Delta u$

Now that we have derived an appropriate loss function for binary classification, we are tasked with the choice of a regularizer term. Much work has been done in the past to find a good regularizer and although there is no clear, best choice we can make do with most suggestions in the literature. In the recent work on this area, Belkin *et al.* have proposed the square of the gradient vector norm, so that:

$$S(u) = \frac{1}{2} \int_{\Omega} \|\nabla u\|^2 d\mathbf{x} , \quad (10)$$

for which  $R(u, \nabla u) = \|\nabla u\|^2$ . The expression of the term depending on  $R$  in Equation 9 is calculated for this choice of  $R$  :

$$\frac{d}{du} \frac{\|\nabla u\|^2}{2} - \nabla \cdot \frac{\partial}{\partial \nabla u} \frac{\|\nabla u\|^2}{2} = -\Delta u. \quad (11)$$

One way to interpret this choice of regularization term is to regard  $u(\mathbf{x})$  as a scalar field and to notice that taking the norm of a vector is interpreted as calculating its magnitude. Thus, this choice minimizes the slope of the gradient points in the direction of the greatest rate of increase. Therefore, it will prevent sharp edges and, in part, will ensure that  $u^*$  be smooth and reduce overfitting.

**Notes.** Equation 10 is convex and thus there is a unique solution  $u^*$  to our problem. This equation is also known as the Dirichlet Energy functional [7], and so we can have a more precise interpretation of minimizing this functional. Chiefly, the Dirichlet Energy measures how variable a function is and it is a quadratic functional on the Sobolev space  $W^{k,2}$ . In general, this regularization term has been tried with some favourable results by Lin *et al.* [8] in the regression context.

## The Completed Model

Without further ado, we shall give the whole model along with its solution. Let  $u : \Omega \subset \mathbb{R}^m \rightarrow \mathbb{R}$  be the function we want to fit and let  $y \in \{0, 1\}$  be the target values. Let  $\sigma(u)$  be the probability that an input  $u = u(\mathbf{x})$  is classified as  $y = 1$ . Let  $F[u]$  be a functional depending on  $u$ , defined by

$$F[u] = \int_{\Omega} -y \ln \sigma(u) - (1 - y) \ln (1 - \sigma(u)) + \|\nabla u\|^2 d\mathbf{x}.$$

Then the condition of optimality is the elliptic PDE

$$\sigma(u) - y - \Delta u = 0 .$$

Here, our equation takes on a more special interpretation. Isolating the variable  $y$  on the left hand side, we see that actual form that  $y = y(\mathbf{x})$  must take is that of the right hand side. Therefore, this expression is the function used to predict the outputs.

$$y(\mathbf{x}) = \sigma(u(\mathbf{x})) - \Delta u(\mathbf{x}) .$$

In order to solve the problem numerically, we can isolate the variable  $y$  to see that we have reduced the original problem to another fitting problem that can be solved numerically. Recalling that  $u = u(\mathbf{x})$  depends on a fixed weight vector  $\mathbf{w}$  and can be expressed by the sum of the set of RBFs  $\{\phi_i(\mathbf{x})\}$ , it is required to find the weight vector  $\mathbf{w}^*$  which appropriately satisfies Equation 12.

$$\underbrace{\sigma(u) - \Delta u}_{\text{Function to fit}} = \underbrace{y}_{\text{Targets}} \quad (12)$$

Since we have  $N$  datum pairs  $(\mathbf{x}_i, y_i)$  we may attempt to utilize these in order to find the function, or equivalently the weights  $\mathbf{w}^*$ , which satisfy the equality above. It is then that we may write the problem of finding  $\mathbf{w}^*$  as a least squares problem (LSQP). Let  $g(\mathbf{x}_i, \mathbf{w}) = \sigma(u(\mathbf{x}_i, \mathbf{w})) - \Delta u(\mathbf{x}_i, \mathbf{w})$ , then the LSQP is:

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_{i=1}^N [y_i + \Delta u(\mathbf{x}_i, \mathbf{w}) - \sigma(u(\mathbf{x}_i, \mathbf{w}))]^2 \quad (13)$$

$$= \arg \min_{\mathbf{w}} \sum_{i=1}^N [y_i - g(\mathbf{x}_i, \mathbf{w})]^2 \quad (14)$$

$$= \arg \min_{\mathbf{w}} \sum_{i=1}^N [y_i - g_i(\mathbf{w})]^2 \quad (15)$$

$$= \arg \min_{\mathbf{w}} S(\mathbf{w}), \quad (16)$$

Note we have defined  $g_i(\mathbf{w}) = g(\mathbf{x}_i, \mathbf{w})$  to obtain the third step. The problem is therefore reduce to solving

this problem. In our particular case, we have chosen to solve this problem via Levenberg-Marquardt [9]; a widely used numerical algorithm used to solve non-linear least squares problems. A detailed description of this algorithm can be found in the book *Numerical Recipes in C* [10]. Using this algorithm will mean that our model will depend on yet another parameter;  $\eta$  is a dampening parameter which will allow to solve a “dampened” version of the problem above.

## Evaluation Methodology

We tested our model against 9 binary datasets. For each dataset we calculated two metrics:

- Accuracy
- Area under the ROC curve (AUC) [11]

Both of the metrics are calculated over the test set in a 5-*fold* cross validation scheme. There has been minimal preprocessing of the datasets. That is, we have centered to 0 mean and unit variance the features of each dataset. Namely, we have *standardized* the dataset. Notably, no dimensionality reduction has been applied.

When training the models, it is necessary to specify the parameters in which they depend on. For our model (LR) we have to choose a parameter triplet  $(c, \lambda, \eta)$ ;  $c$  is the fitting degree,  $\lambda$  is the regularization parameter and  $\eta$  is a dampening parameter. For a *RBF-kernel* SVM, we choose the parameter pair  $(C, \gamma)$ ;  $C$  is the penalty parameter of the error term while  $\gamma$  is the kernel coefficient. Lastly, we assume that NN is a *multi-layer perceptron* with one hidden layer of 100 nodes which depends on a regularization term  $\alpha$ . Both SVM and NN are implementations of the Python library `sklearn` [12]. The parameters for each of our models were chosen by the following methodology:

**LR.** The parameter triplet  $(c, \lambda, \eta)$  is searched on a grid:

- $c$  is searched on the interval  $(0, 5)$  with step =  $\ln 2$ .
- $\lambda$  is searched on the interval  $[0, 10]$  with step = 1.
- $\eta$  is fixed to  $\eta = 1$  for each and every dataset.

The AUC score may also be used to further determine the performance of a classifier [13]. Proceeding in a similar manner as before, we calculate the average residual between the 1st place and LR. We find that LR is on average down from 0.0126 units from the top performer for each dataset. The same is calculated for SVM and

The value of  $\eta = 1$  was found empirically to work well with almost any value of  $\lambda$  and  $c$ . To threshold the outputs of LR into values of the set  $\{0, 1\}$  we will use the following classification rule:

$$g(\mathbf{x}) = \begin{cases} 0 & \sigma(u(\mathbf{x})) - \Delta u(\mathbf{x}) \leq \frac{1}{2} \\ 1 & \sigma(u(\mathbf{x})) - \Delta u(\mathbf{x}) > \frac{1}{2} \end{cases}$$

**SVM.** The parameter pair  $(C, \gamma)$  is searched on a grid:

- $C$  is searched on the interval  $(0, 5)$  with step =  $\ln 2$ .
- $\gamma$  is fixed to  $\gamma = 1/m$ , where  $m$  is the number of features of the dataset.

Originally  $\gamma$  was searched on the interval  $[0, 10]$  with step = 1. It was soon found empirically, that fixing  $\gamma$  to  $\gamma = 1/m$  resulted in better performance.

**NN.** The parameter  $\alpha$  was fixed to  $\alpha = 0.001$ .

## Results and Analysis

This section presents the results in Tables 1 through 3. To understand more easily the performance of the methods, we have arranged the results into various tables. In terms of accuracy, LR outperformed SVM and NN on 3 datasets. In detail, LR outperformed NN on 4 datasets and outperformed SVM on 2, tying on the Breast Cancer dataset. To fully grasp how much better or worse our method has performed we calculated the absolute value of the residual between LR and the top performer for each dataset. On average we see that our method was down by 0.9567 %. Although not shown in Table 1, the average distance from the top performer for SVM is 0.5229 % and for NN it is 0.8975 %. These scores lend to the interpretation of “which method got closer to the real solution”. Looking at who got the top score we might say that the top performer was the best *method suited for that particular type of dataset*. It is of interest then to evaluate a method which in average should perform well for most types of datasets. Interpreting the results, we can confidently say that **SVM is the best method, while NN is second and LR comes a close third in terms of accuracy.**

NN; respectively 0.0088 and 0.0161. Surprisingly, even though NN outperformed SVM and LR on 4 different datasets while LR only outperformed the others on one, **on average LR will perform better than NN.** Unsurprisingly, SVM will still perform better than NN and LR.

Table 1: The accuracy (%) of each method is outlined in this table. The second to last column indicates which ranking (1st, 2nd or 3rd) LR obtained; higher is better. On the other hand, the last column indicates the absolute value of the residual between LR and the 1st place. In bold, the best accuracy.

Data	Dim	$N$	LR	SVM	NN	Place	Dist. from 1st
Australian	14	690	86.6667	86.8116	<b>87.8261</b>	3rd	1.1594
Blood Transfusion	4	748	<b>78.2246</b>	78.2237	77.2859	1st	0.0
Breast Cancer	30	569	97.7146	<b>98.7425</b>	98.2673	3rd	1.0279
Bupa	6	345	<b>72.4638</b>	<b>72.4638</b>	71.0145	1st	0.0
German	24	1000	76.0000	76.6000	<b>78.3000</b>	3rd	2.3
Haberman	3	306	73.5431	73.8710	<b>74.5267</b>	3rd	0.9836
Heart	13	270	82.2222	<b>84.8741</b>	84.0148	3rd	2.6519
Sonar	60	208	88.4321	<b>88.9199</b>	87.4681	2nd	0.4878
Vertebral Column	6	310	<b>86.7742</b>	85.4839	83.8710	1st	0.0
Average distance from 1st:							0.9567

Table 2: The area under the ROC curve has been calculated for each method over each dataset Also, we have added a grade next to the scores to see how they perform against each other more easily; A being excellent performance, while F is catalogued as a fail. Likewise (Table 1), the second to last column indicates which ranking (1st, 2nd or 3rd) LR obtained; higher is better. On the other hand, the last column indicates the absolute value of the residual between LR and the 1st place. In bold, the best AUC.

Data	LR	SVM	NN	Place	Dist. from 1st
Australian	0.8643 (B)	0.8701(B)	<b>0.8777</b> (B)	3rd	0.0134
Blood Transfusion	0.5844 (F)	<b>0.6144</b> (D)	0.5502 (F)	2nd	0.03
Breast Cancer	0.9729 (A)	0.9800 (A)	<b>0.9858</b> (A)	3rd	0.0129
Bupa	0.7024 (C)	<b>0.7059</b> (C)	0.6866 (D)	2nd	0.0035
German	0.6821 (D)	0.6920 (D)	<b>0.7058</b> (C)	3rd	0.0237
Haberman	<b>0.5560</b> (F)	0.5559 (F)	0.5463 (F)	1st	0.0
Heart	0.8180 (B)	0.8452 (B)	<b>0.8322</b> (B)	3rd	0.0142
Sonar	0.8857 (B)	<b>0.8906</b> (B)	0.8812 (B)	2nd	0.0049
Vertebral Column	0.8292 (B)	<b>0.8404</b> (B)	0.7978 (C)	2nd	0.0112
Average distance from 1st:					0.0126

Even though our analysis has been very exact until now, it is time to present a more intuitive analysis based on the AUC. Furthermore, this analysis is more *robust* than the previous. The analysis consists in assigning a “grade” to a classifier by specifying the following grading scheme:

$$\text{Grade} = \begin{cases} \text{Excellent (A)} & 0.9 \leq \text{AUC} \leq 1 \\ \text{Good (B)} & 0.8 \leq \text{AUC} < 0.9 \\ \text{Fair (C)} & 0.7 \leq \text{AUC} < 0.8 \\ \text{Poor (D)} & 0.6 \leq \text{AUC} < 0.7 \\ \text{Fail (F)} & 0.5 \leq \text{AUC} < 0.6 \end{cases}$$

The “robustness” comes from the fact that we are partitioning discretely the interval  $[0, 1]$  and assigning each a grade. Small variations within the sub-intervals will be neglected. Looking at Table 2 we see that each of the

AUC scores has a letter assigned to it. This is interpreted as the grade which the method received on that particular dataset. In order to summarize the grades, we arrange the number of times a method received a particular grade in Table 3. Simply by looking at Tables 2 and 3 we can get the sense that all of the models performed similarly. In order to obtain a quantitative measure, we may assign each grade a value. Namely,  $A = 1, B = 2, \dots, F = 5$ . Obtaining the weighted total of the grades will then let us asses directly which method is better by looking at the lowest total. The results are presented in Table 3. Immediately, we see that SVM once again obtained the best score. This time, however, LR came a close second with only a 1 point difference. On the contrary NN was down by 4 points from SVM and 3 points from LR. Summarizing, **LR outperforms NN again while SVM remains overall the best method.**

Table 3: The columns LR, SVM and NN indicate the numbers of times each of the methods got the grade on the leftmost column. Letting A = 1, B = 2, C = 3, D = 4, F = 5, we can calculate a weighted final grade for our classifiers and see how each of them performed. Clearly a lower grade is better. The weighted grade is just the weighted total of the grades each of these methods obtained.

Grade	LR	SVM	NN
A	1	1	1
B	4	4	3
C	1	1	2
D	1	2	1
F	2	1	2
Weighted grade (lower is better):	26	<b>25</b>	29

## Conclusion

Even though supervised learning has had tremendous advances in the last few years, it remains clear that a lot of work has yet to be done. As part of our contribution to the field, we have designed and implemented a new variational model for binary classification. Opting to attack the supervised learning problem by functional and variational means, we arrived at the LR model and gave each of its components interpretations. Moreover we used the Levenberg-Marquardt method to solve the optimality condition which in essence was equivalent to “training the model”. We outlined our evaluation criteria and methodology so that we could compare how well our method stood against two of the most utilized methods; NN and SVM. From the results obtained, we conducted three different analyses to quantify the performance of each model. The accuracy results obtained, clearly indicated that SVM was the superior method amongst the three, and so the focus shifted to the com-

parison between NN and LR. It was found that NN outperformed LR by a small margin. The second analysis showed that LR outperformed NN while SVM came out on top once again. Finally, we conducted a more robust analysis which permitted us to ignore small variations in the scores. This last analysis further showed that LR clearly outperformed NN. Moreover it showed that LR was not far from SVM’s performance. We speculate that this is due to the heavy optimizations in the implementation of SVM which permit it to find the minimum in a faster manner.

Without a doubt, the variational approach looks to be very promising and is to be explored in future work. Namely, we will concentrate in deriving expressions for higher order derivatives and better formalizing the problem at hand in order to derive more interesting properties. The implementations for the LR model, benchmarking scripts and datasets may all be found at <https://github.com/carlosb/thesis>.

## References

- [1] Goodfellow, Ian, Yoshua Bengio, and Aaron Courville: *Deep learning*. MIT press, 2016.
- [2] A. Singh, N. Thakur and A. Sharma, A review of supervised machine learning algorithms, 2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom), New Delhi, 2016, pp. 1310-1315.
- [3] Fern’andez-Delgado, Manuel, et al., Do we need hundreds of classifiers to solve real world classification problems? *J. Mach. Learn. Res* 15.1 (2014): 3133-3181.
- [4] Kotsiantis, Sotiris B., I. Zaharakis, and P. Pintelas: Supervised machine learning: A review of classification techniques. (2007): 3-24.
- [5] Caruana, Rich, and Alexandru Niculescu-Mizil, An empirical comparison of supervised learning algorithms. Proceedings of the 23rd international conference on Machine learning. ACM, 2006.
- [6] Shore, John, and Rodney Johnson, Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on information theory* 26.1 (1980): 26-37.
- [7] Evans, Lawrence. Partial differential equations. (2010).

- [8] Lin, Tong, et al., Supervised learning via Euler’s Elastica models, *Journal of Machine Learning Research* 16 (2015): 3637-3686.
- [9] Mor’e, Jorge J. *The Levenberg-Marquardt algorithm: implementation and theory*, Numerical analysis. Springer, Berlin, Heidelberg, 1978. 105-116.
- [10] Press, William H., et al., *Numerical recipes in C. Vol. 2*. Cambridge: Cambridge university press, 1996.
- [11] Hanley, James A., and Barbara J. McNeil, The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* 143.1 (1982): 29-36.
- [12] Pedregosa, Fabian, et al., Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12.Oct (2011): 2825-2830.
- [13] Fawcett, Tom, An introduction to ROC analysis. *Pattern recognition letters* 27.8 (2006): 861-874.
- [14] Belkin, Mikhail, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research* 7.Nov (2006): 2399-2434.