Title	CascadeML: An Automatic Neural Network Architecture Evolution and Training Algorithm for Multi-label Classification
Authors(s)	Pakrashi, Arjun, MacNamee, Brian
Publication date	2019-11-19
Publication information	Pakrashi, Arjun, and Brian MacNamee. "CascadeML: An Automatic Neural Network Architecture Evolution and Training Algorithm for Multi-Label Classification," 2019.
Conference details	The 39th SGAI International Conference on Artificial Intelligence (AI 2019), Cambridge, United Kingdom, 17-19 December 2019
Series	Lecture Notes in Computer Science, 11927, Lecture Notes in Artificial Intelligence, 11927
Item record/more information	http://hdl.handle.net/10197/11677
Publisher's statement	The final publication is available at www.springerlink.com .
Publisher's version (DOI)	10.1007/978-3-030-34885-4_1

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/332630595>

CascadeML: An Automatic Neural Network Architecture Evolution and Training Algorithm for Multi-label Classification

Preprint · April 2019

CITATIONS

0

READS

34

2 authors:

[Arjun Pakrashi](#)

University College Dublin

10 PUBLICATIONS 19 CITATIONS

[SEE PROFILE](#)

[Brian Mac Namee](#)

University College Dublin

103 PUBLICATIONS 617 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:

Active Learning for Regression [View project](#)

CascadeML: An Automatic Neural Network Architecture Evolution and Training Algorithm for Multi-label Classification^{*}

Arjun Pakrashi and Brian Mac Namee

Insight Centre for Data Analytics
University College Dublin
Dublin, Ireland

arjun.pakrashi@ucdconnect.ie, arjun.pakrashi@insight-centre.org
brian.macnamee@ucd.ie

Abstract. Multi-label classification is an approach which allows a datapoint to be labelled with more than one class at the same time. A common but trivial approach is to train individual binary classifiers per label, but the performance can be improved by considering associations within the labels. Like with any machine learning algorithm, hyperparameter tuning is important to train a good multi-label classifier model. The task of selecting the best hyperparameter settings for an algorithm is an optimisation problem. Very limited work has been done on automatic hyperparameter tuning and AutoML in the multi-label domain. This paper attempts to fill this gap by proposing a neural network algorithm, CascadeML, to train multi-label neural network based on cascade neural networks. This method requires minimal or no hyperparameter tuning and also considers pairwise label associations. The cascade algorithm grows the network architecture incrementally in a two phase process as it learns the weights using adaptive first order gradient algorithm, therefore omitting the requirement of preselecting the number of hidden layers, nodes and the learning rate. The method was tested on 10 multi-label datasets and compared with other multi-label classification algorithms. Results show that CascadeML performs very well without hyperparameter tuning.

1 Introduction

In *multi-label classification* problems a datapoint can be assigned to more than one class, or *label*, simultaneously [12]. For example, an image can be classified as containing multiple different objects, or music can be labelled with more than one genre. This contrasts with *multi-class classification* problems in which objects can only belong to a single class. Multi-label classification algorithms either break the multi-label problem down into smaller multi-class classification problems—for example *classifier chains* [23]—and are known as *problem transformation* methods—or modify multi-class algorithms to directly train

^{*} This research was supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.

on multi-label datasets—for example *BackPropagation in Multi-Label Learning* (BPMLL) [37]—and are known as *algorithm adaptation* methods.

Automatic machine learning [8], or *AutoML*, approaches have seen a recent resurgence of interest as researchers look for ways to automatically select optimal algorithms, features, model architectures, and hyperparameters for machine learning tasks. The AutoML research community has, however, paid very little attention to multi-label classification problems, although there have been some recent efforts [25, 26, 33].

The *Cascade2* algorithm [21] is an interesting neural network approach that learns model parameters and model architecture at the same time. In *Cascade2*, which is based on the cascade correlation neural network approach [7], training starts with a simple perceptron network, which is grown incrementally by adding new cascaded layers with skip-level connections as long as performance on a validation dataset improves. Weights in each new layer are trained independently of the overall network which greatly reduces the processing burden of this approach.

This paper proposes *CascadeML*, a new AutoML solution for multi-label classification problems, that is inspired by the *Cascade2* algorithm and BPMLL. Improvements are made to both components leading to an implementation that requires minimal hyperparameter or network architecture tuning. In a series of evaluation experiments this approach has been shown to perform very well without the extensive hyperparameter tuning required by state-of-the-art multi-label classification methods. To the best of authors’ knowledge this is the first automatic neural network architecture selection and training approach for multi-label classification methods.

The remainder of the paper is structured as follows. Section 2 discusses the existing literature including a formal definition of multi-label classification and the BPMLL algorithm. Section 2.3 describes the cascade neural network approach and, specifically, the *Cascade2* algorithm. The proposed *CascadeML* method is then presented in Section 3. The design of an experiment to evaluate the performance of the *CascadeML* algorithm, and benchmarking its performance against state-of-the-art multi-label classification approaches is described in Section 4. Section 5 presents and discusses the results of this experiment. Finally, Section 6 discusses future research directions and concludes the paper.

2 Related Work

In this section first the cost function of BPMLL will be mentioned followed by a brief review of AutoML in multi-label literature. Then the *Cascade2* algorithm will be explained.

2.1 BPMLL

The first neural-network based multi-label algorithm, *BackPropagation in Multi-Label Learning* (BPMLL), was proposed by Zhang et al. in 2006 [37]. It is a single

hidden layer, fully connected feed-forward architecture, which uses the back-propagation of error algorithm to optimise a variation of the ranking loss function [38] that takes pairwise label associations into account. This loss function can be defined as follows:

$$E = \sum_{i=1}^n \frac{1}{|\mathbf{y}_i||\bar{\mathbf{y}}_i|} \sum_{(k,l) \in (\mathbf{y}_i \times \bar{\mathbf{y}}_i)} \exp(-(c_i^{(k)} - c_i^{(l)})) \quad (1)$$

Here $\mathbf{y}_i$ indicates the set of labels assigned to $\mathbf{x}_i$ and $\bar{\mathbf{y}}_i$ indicates the set of labels which are not assigned to $\mathbf{x}_i$. The network uses the *tanh* activation function, therefore this algorithm uses a bipolar encoding of the target variables: $y_i^{(l)} = +1$ if the label l is relevant to $\mathbf{x}_i$, and -1 if irrelevant. Here $c_i^{(k)}$ and $c_i^{(l)}$ are the outputs of the k^{th} and the l^{th} output units representing the corresponding label predictions for the datapoint $\mathbf{x}_i$.

The intuition behind this loss function is that for a pair of labels (k, l) , where k is relevant to the datapoint $\mathbf{x}_i$ and l is not, if the prediction score for k is positive whereas the prediction score for l is negative, then $\exp(-(c_i^{(k)} - c_i^{(l)}))$ has the minimum penalty. An incorrect prediction score order results in higher penalty. Therefore, minimising Eq. (1) would result in pairs of labels being predicted correctly.

For BPMLL, like any neural network algorithm, the number of hidden units has to be determined, which is a hyperparameter to be tuned. In [10] modifications to the BPMLL loss function were proposed. This modified version learns the network as in BPMLL, and also learns the values using which the predicted scores are thresholded to get label assignments.

There have been a small number of other neural network approaches specifically designed for multi-label classification scenarios. In 2009, Zhang et al. proposed a *multi-label-based radial basis function* network, ML-RBF [36]. This is an extension of the RBF network, optimising the sum-of-squares function. *Multi-class multi-label perceptron* (MMP) [5] trains perceptrons for each label but in a way such that the applicable labels are ranked higher than the incorrect labels, thus considering associations between labels. An improvement of MMP, *multi-label pairwise perceptron* (MLPP), was proposed in [16]. This approach trains the perceptrons for each pair of classes. Nam, et al. [17] demonstrate the efficiency and effectiveness of cross-entropy for multi-label classification, improving the work of BPMLL by using several recent developments such as ReLU activation units, dropout and the use of the adaptive gradient descent algorithm AdaGrad [9].

Some work involving deep neural networks on computer vision and image recognition were done in [3, 34, 39, 40], which uses multi-label datasets as a part of the training pipeline. Similarly, convolutional neural networks was extended to predict multi-label images in [32]. In [22] the feature space of multi-label classification was modified using deep belief networks such that the labels become less dependent, after which well-known multi-label algorithms are applied in the modified space.

2.2 AutoML

AutoML algorithms focusing on multi-label specific problems are approached in [25, 26], using genetic algorithms to train and select multi-label models. Wever et al. [33] propose an extension of an existing multi-class AutoML tool for multi-label. Except these works, no other AutoML based or automatic hyperparameter tuning based work on the multi-label domain was found.

The cascade correlation neural network approach [7] was an early AutoML method. In cascade correlation neural networks training starts with a simple perceptron network, which is grown incrementally by adding new cascaded layers with skip-level connections as long as performance on a validation dataset improves. Since the proposal of the original cascade correlation algorithm in [7], various improvements that follow a similar overall process to the original method have been proposed, for example in [1, 11, 20, 31], as well as Cascade2 [21]. Active research in this field, however, is fairly limited. As it is the basis for CascadeML, the Cascade2 algorithm is described in detail in the next section.

2.3 The Cascade2 Algorithm

This section describes the Cascade2 algorithm upon which CascadeML is based. The generic architecture of a cascade neural network is first described, before the specific Cascade2 training algorithm is presented.

Architecture The *cascade correlation neural network*, first proposed by Fahlman & Lebiere [7], is an incremental greedy multi-class neural network learning algorithm which grows the network architecture at the same time as it trains the network weights. For a multi-class classification problem with d inputs and q classes, the architecture of a network trained using CascadeML will have $d + 1$ inputs (including a bias term) and q outputs (one for each class). Each of the network's L hidden layers, l_i , will have only one unit, which receives incoming weights from all the $d + 1$ inputs as well as from all the hidden units in the previous layers. The output of each hidden layer l_i is connected to the q outputs of the network. A layer with such a connection scheme is called a *cascade layer*.

We can categorise the weights in a cascade network into four types:

1. **Input to output layer weights** connecting the $d + 1$ inputs to the q outputs, forming a perceptron network.
2. **Input to hidden layer weights** connecting the $d + 1$ inputs to the L hidden cascade layers.
3. **Hidden to hidden layer weights** connecting the output of all the previous hidden cascade layers $l_1, l_2, \dots, l_{i-1}$, to the hidden cascade layer l_i .
4. **Hidden to output layer weights** connecting the outputs of the cascade layers $l_1, l_2, \dots, l_L$ to the q output units.

Figure 1h shows an example of a simple cascade neural network with three inputs, two output classes, and three hidden cascade layers (l_1 , l_2 , and l_3). All

Fig. 1: The steps in the Cascade 2 network training algorithm. In each diagram the circles labelled *Inputs* correspond to the input layer and the circles labelled *Outputs* correspond to the output layer of the network. Hidden cascade units are represented by the circles labelled l_i . A weight between nodes in two layers exists, where horizontal and vertical lines intersect. Crosses indicate a weight that is trainable at a specific step in the training process, while squares indicate a weight that is frozen at a specific step. g is the activation function used at the output layer. (h) shows a more typical network diagram of the final network trained.

connections flow from left to right. The cascade network is grown dynamically, one layer at a time, and the four different types of weights are each trained in slightly different ways (explained in detail below). Once training is complete, prediction uses a straight-forward feedforward algorithm that propagates values through the cascade layers.

Training Model training in the Cascade2 algorithm starts with a simple perceptron network (Figure 1a) with $d + 1$ inputs and q outputs. This network is

referred to as the *main network*. The main network is grown as training proceeds by iteratively adding hidden cascade layers to it. This is achieved by iteratively repeating two phases, *Phase I* and *Phase II*, each of which trains different parts of the cascade network.

In Phase I, the input to output layer weights (type 1 in the list above) and hidden to output layer weights (type 4) of the main network are trained, while all other weights (input to hidden layer and hidden to hidden layer) are frozen. The target values used in this phase to calculate the loss of the network are the target classes from the original dataset. The mean squared error (MSE) between the output of the main network and the ground truth is minimised using gradient descent.

Phase II trains and adds a new cascade layer l_i at the i^{th} iteration of training. The inputs to the newly added layer l_i are the $d + 1$ input dimensions, and the outputs from the previous hidden layers $l_1, \dots, l_{i-1}$, in the main network. At this phase only the weights involving the new hidden layer, l_i , are trained. These are the input to hidden layer weights (type 2) for l_i ; hidden to hidden layer weights on connections of the output of previous hidden cascade layers, $l_1, \dots, l_{i-1}$, to the current hidden layer, l_i (type 3); and the weights connecting the new hidden layer, l_i , to the output layer (type 4). All other weights in the main network are frozen. In this phase the target values used in training are not the original target values, but rather the error between the MSE of the main network constructed up to the previous iteration $i - 1$, and the output of the new layer l_i .

Once the weights associated with the new hidden layer have been trained the layer and these weights are added to the main network. The weights connecting the new hidden layer, l_i , to the output layer are negated when these are added to the main network. This is so that the contribution of the output of the newly added layer will minimise the error of the main network [18]—recall that the newly added layer was trained to predict the main network error.

When Phase II is complete, the algorithm proceeds again to Phase I and continues iterating between Phase I and Phase II until a maximum depth is reached or a learning error threshold is not exceeded. Training always ends with Phase I.

Example Figure 1 shows an example of the growth of a cascade network (the neural network diagram scheme used by Fahlman & Lebiere [7] is used). Figure 1a shows the initial network with 3 inputs and 2 outputs. In this schematic the intersections of the straight lines indicate the weights. A cross at an intersection indicates that a weight is trainable at the current phase, while a square indicates that a weight is frozen. The algorithm starts in Phase I and the network in Figure 1a is trained. All input to output layer weights (type 1), are trained (no hidden to output layer weights (type 4) exist yet). Next, in Phase II, a new cascade layer, l_1 , is added as shown in Figure 1b, and only the input to hidden layer (type 2) and hidden to hidden layer (type 3) weights related to the newly added layer, l_1 , are trained. Next, the process goes back to Phase I and trains input to output layer (type 1) and hidden to output layer (type 4) weights in the main

network as shown in Figure 1c. This process iterates two more times through Phase I and II in Figures 1d, 1e and 1f until the final network in Figure 1g is produced. Figure 1h shows this same final network using a more typical network diagram.

3 The CascadeML Algorithm

CascadeML is a cascaded neural network approach to multi-label classification based on Cascade2 [21]. The main objective of this method is to find good multi-label classifier models that take advantage of label associations, while minimising the model selection and training time by omitting hyperparameter tuning and architecture tuning.

CascadeML uses a similar training process to that described in Section 2.3. CascadeML starts with a perception network with $d + 1$ inputs (including the bias unit) and q output units, one for each label. In Phase I, only the hidden to output layer and input to output layer weights are trained, as in Cascade2. The loss function used in this phase is the BPMLL loss function shown in Eq. (1), which allows CascadeML to consider pairwise label associations.

In Phase II CascadeML differs from Cascade2 in the following way. First, instead of adding hidden cascade layers with a single unit at each iteration, hidden cascade layers with multiple units are added. This gives rise to a hyperparameter selection problem as the number of units in each hidden layer needs to be determined. To overcome this, at each iteration of CascadeML, a *candidate pool* of many candidate hidden cascade layers is trained, that could be added to the main network. Each of the candidate hidden cascade layers is initialised with randomly selected initial weight values, a randomly selected activation function, and a randomly selected number of units. Each of the candidate hidden cascade layers is trained independently in parallel, to minimise MSE as explained in the Cascade2 algorithm. Once they have all been trained the best candidate hidden cascade layer from the candidate pool is selected (based on calculated loss on a validation dataset) and added to the main network.

To add flexibility to the network architectures explored by Cascade-ML, the algorithm can include candidate hidden cascade layers that are *sibling* layers to the deepest hidden cascade layer already in the main network [1], as well as *successor* cascade layers. This allows wide architectures as well as deep architectures to be explored. This is done by training candidate cascade networks in the candidate pool as successors and siblings and then selecting the best of the two types of candidate network.

The candidate hidden cascade layers in the candidate pool can each be trained independently in isolation from the main network, because when training the candidate hidden cascade layer, l_i , the inputs to the layer, the targets and the weights of the main network are all fixed. Therefore, the hidden cascade layer, l_i , can be considered a subnetwork, trained in isolation and then added to the main network.

Table 1: Multi-label datasets

Dataset	Instances	Inputs	Labels	Labelsets	Cardinality	MeanIR
flags	194	26	7	24	3.392	2.255
yeast	2417	103	14	77	4.237	7.197
scene	2407	294	6	3	1.074	1.254
emotions	593	72	6	4	1.869	1.478
medical	978	1449	45	33	1.245	89.501
enron	1702	1001	53	573	3.378	73.953
birds	322	260	20	55	1.503	13.004
genbase	662	1186	27	10	1.252	37.315
cal500	502	68	174	502	26.044	20.578
llog	1460	1004	75	189	1.180	39.267

When the best candidate hidden cascade layer is selected from the candidate pool, it is added to the main network by copying the input to hidden layer, weights to the main network, negating the hidden to output layer weights and connecting them to the main network as in Cascade2. The main network increases in depth or the deepest layer grows in breadth depending on whether a successor or a sibling candidate layer was selected.

For both Phase I and Phase II, an adaptive first order gradient descent algorithm iRProp- [13], a variant of RProp [13, 24], is used. iRProp- was found to be more stable than the originally used Quickprop [6]. iRProp- is an adaptive algorithm which uses an adaptive learning rate and the sign of the partial derivative of the error function for each weight adjustment. This method mainly helps learn very fast in the flat regions of the error space and near local minima, as it uses only the sign of the partial derivative (ignoring its magnitude) and uses an adaptive learning rate. L2 regularisation [9] was used in both phases of CascadeML.

4 Experiment Design

To evaluate the effectiveness of CascadeML, an experiment was performed on ten well-known multi-label benchmark datasets listed in Table 1. In Table 1 *Instances*, *Inputs* and *Labels* are the number of datapoints, the dimension of the datapoint and the number of labels respectively. *Labelsets* indicates the number of unique combinations of labels present in a dataset. *Cardinality* measures the average number of labels assigned to each datapoint, and *MeanIR* [2] indicates the imbalance ratio of the labels.

The performance of models trained using CascadeML was compared with the multi-label neural network algorithm BPMLL, and four other state-of-the-art multi-label classification algorithms: *classifier chains* [23] (CC), *RAkEL* [30], HOMER [28], and MLkNN [35]. These algorithms were selected to cover different types of multi-label classification techniques. Classifier chains, RAKEL and HOMER, when used with SVMs, are ensemble methods that have been previously shown to be the best performing the multi-label classifiers [15, 19]; BPMLL

is a well-known multi-label specific neural network algorithm; and MLkNN is a nearest-neighbour based algorithm adaptation method. The implementations of classifier chains, RAkEL, HOMER, MLkNN and BPMLL are from the MULAN library [29] and implemented in Java. CascadeML was implemented in Python¹.

To compare the performances of the methods, label-based macro-averaged F-Score [38] was used. This is preferred over Hamming loss [38], used in several previous studies (e.g. [4,27,35]), as when used with highly imbalanced multi-label datasets Hamming loss tends to allow the performance on the majority labels to overwhelm performance on the minority labels. Label-based macro-averaged F-Score does not suffer from this problem. For every dataset performance is evaluated using a 2 times 5-fold cross validation experiment. The mean label-based Macro-averaged F-Score from these experiments are reported.

4.1 Configuring CascadeML

Although there is no hyperparameter tuning required for CascadeML, it does require some configuration. In the experiments described here, at each iteration, the candidate pool contained two candidates for each combination of layer type—*successor* or *sibling*—and activation unit type—*linear*, *sigmoid*, or *tanh*. This made for 12 candidate hidden cascade layers at each iteration. The number of hidden units in each candidate layer was randomly selected as a fraction of the number of input dimensions, d , following a uniform distribution in $(0, 1]$.

For the output layer of the main network the activation function used was *tanh* as the cost function requires bipolar encoding of the labels. During Phase II of training the outputs of the candidate layers in the pool use a linear activation function, as explained in Section 3. L2 regularisation was used in all training phases with regularisation value of 10^{-5} . In Phase I and Phase II training early stopping is used where training stops if the average loss (based on a validation dataset) calculated over a window of the last 20 training epochs increases from one iteration to the next. For both Phase I and Phase II iRProp- is initialised as recommended in [13]. The maximum number of iterations allowed for iRProp- in both phases was 2000. To set an upper bound on network growth in CascadeML two stopping criteria were used: (1) a new cascade layer (sibling or successor) was added only if did not lead to an increase in the validation loss of the entire network, and (2) only 20 iterations are allowed.

4.2 Configuring other algorithms

All the algorithms used in the experiment, except CascadeML, underwent a grid search based hyperparameter tuning using 2 time 5-folds cross validation. For classifier chains, RAkEL and HOMER, support vector machines [14] with a radial basis kernel (SVM-RBF) were used as the base classifier. In these cases 12 combinations of the regularisation parameter, C , and the kernel spread, σ , were included the hyperparameter grid. For RAkEL the subset size hyperparameter

¹ A version of CascadeML is available at: <https://github.com/phoxis/CascadeML>

Table 2: Results of experiments. Rows indicate the datasets, columns indicate algorithms. Values in cells are mean label-based macro-averaged F-Scores and the standard deviations followed by relative ranks in parenthesis. Last row are the average ranks of the corresponding algorithms.

	CascadeML	RAkEL	CC	BPMLL	HOMER	MLkNN
flags	0.6723±0.06 (1)	0.6505±0.04 (2)	0.6405±0.06 (4)	0.5948±0.03 (6)	0.6479±0.04 (3)	0.6009±0.07 (5)
yeast	0.4624±0.01 (1)	0.4367±0.02 (4)	0.4510±0.01 (2)	0.4357±0.01 (5)	0.4478±0.02 (3)	0.3772±0.01 (6)
scene	0.7606±0.01 (5)	0.8017±0.01 (2)	0.8040±0.01 (1)	0.7777±0.01 (4)	0.8001±0.02 (3)	0.7424±0.02 (6)
emotions	0.6671±0.02 (2)	0.6281±0.02 (4)	0.6242±0.01 (5)	0.6899±0.02 (1)	0.6212±0.02 (6)	0.6294±0.03 (3)
medical	0.6758±0.02 (3)	0.6966±0.03 (1)	0.6924±0.04 (2)	0.5582±0.08 (5)	0.6108±0.05 (4)	0.5398±0.05 (6)
enron	0.2852±0.02 (3)	0.2882±0.04 (2)	0.2890±0.03 (1)	0.2806±0.02 (5)	0.2812±0.03 (4)	0.1771±0.03 (6)
birds	0.4812±0.03 (1)	0.1812±0.06 (4)	0.1582±0.06 (5)	0.3426±0.06 (2)	0.1551±0.05 (6)	0.2256±0.09 (3)
genbase	0.9403±0.02 (3)	0.9432±0.05 (2)	0.9440±0.04 (1)	0.8149±0.12 (6)	0.9394±0.05 (4)	0.8502±0.05 (5)
cal500	0.2263±0.01 (2)	0.1790±0.01 (5)	0.1849±0.01 (4)	0.2367±0.02 (1)	0.1988±0.02 (3)	0.1007±0.01 (6)
llog	0.2264±0.03 (6)	0.2998±0.05 (1)	0.2916±0.03 (3)	0.2953±0.06 (2)	0.2561±0.03 (5)	0.2630±0.05 (4)
Avg. rank	2.7	2.7	2.8	3.7	4.1	5.0

(ranging from 2 to 6) was also included, and for HOMER the cluster size hyperparameter (ranging from 2 to 6) was also included. For BPMLL the only hyperparameter in the grid search was the number of units in the hidden layer. Sizes of 20%, 40%, 60%, 80% and 100% of the number of inputs for each dataset were explored, as recommended by Zhang et al. [37]. In this case the L2 regularisation coefficient was set to 10^{-5} and a maximum of 10000 iterations were allowed, based on [19].

The results presented are based on the best performing hyperparameter combinations. Finally, the mean label-based Macro-averages F-Scores of 2 times 5-folds cross validation experiments of the best hyperparameter combination are reported.

5 Results

The results of the experiments are shown in Table 2, where the columns indicate the algorithms and the rows indicate the datasets. Each cell of the table shows the label-based macro-averaged F-Score (higher values are better) followed by the standard deviation over the cross validation folds. These label-based F-Scores are computed through extensive cross validated hyperparameter tuning. The values in the parenthesis indicate the relative ranking (lower values are better) of the algorithm with respect to the corresponding dataset. The last row of Table 2 indicates the overall average ranks of the corresponding algorithms.

Table 2 shows that CascadeML (avg. rank 2.7) performed better than BPMLL (avg. rank 3.7), HOMER (avg. rank 4.1), MLkNN (avg. rank 5.0) and CC (avg. rank 2.8), overall. RAkEL had the same overall average rank as CascadeML.

Although RAkEL and CC had similar rank as CascadeML on average, it must be noted that the label-based macro-averaged F-Score for CC as well as for the other methods were achieved after doing an extensive hyperparameter tuning which CascadeML did not require. Besides tuning for C and sigma hyperparameter of the underlying SVM-RBFs for RAkEL, CC and HOMER, there are other hyperparameters which needs to be tuned. For RAkEL the labelset size

Fig. 2: Histogram of label-based macro-averaged F-Scores achieved from all hyperparameter combinations of subset size for RAKEL and C, sigma for the underlying SVM-RBFs. The vertical dotted line indicates CascadeML’s performance.

needs to be selected, for CC the chain order needs to be defined, and for HOMER the clustering algorithm and the cluster size needs to be defined. All of these hyperparameter increases the hyperparameter search space dimensionality.

Absolute running times or the number of operations are not directly comparable as the methods are different from CascadeML and implementations of the algorithms span different programming languages. However, it is worth noting that the completion of the CC, RAKEL benchmarks took multiple weeks (with multiple folds run in parallel) due to the hyperparameter tuning involved, whereas running the equivalent benchmark for CascadeML took less than a week.

The nature of the incremental growth and training in combination with the fast convergence nature of iRProp- with the L2 regularisation helped the network to generalise as well as converge faster. Also, note that the candidate unit pool size was set to 12 and all of them were run in parallel, hence the real runtime of the candidate training would be the maximum time taken of the 12 of the candidates. Therefore, by exploiting the cascade architecture and training process, as well as using the iRProp- algorithm along with L2 regularisation, CascadeML was able to maintain a very good level of performance without hyperparameter tuning.

Figure 3d shows the training costs for scene dataset for one fold. The vertical dotted line indicates the addition of a candidate layer to the main network. After each addition of the candidate network the cost increases but then sharply decreases at first then continues decreasing steadily.

It is important to note that CascadeML has the advantage of not requiring hyperparameter tuning. For other algorithms the selection of hyperparameter values can have a huge impact on performance. For example, Figure 2 shows the distribution of the label-based macro-averaged F-Scores for different combinations of label subset size, C and sigma values the underlying SVM-RBFs of RAKEL for the yeast and enron datasets. Note that the F-Score values in Figure

Fig. 3: CascadeML trained network properties.

2 vary significantly. For the yeast dataset, CascadeML performed the best, and for enron dataset only 2.1 % of the hyperparameter attained better result than CascadeML. In general the distribution skews towards models with relatively poorer performance. CascadeML attained similar high values of performance in both the cases (0.4624 for the yeast dataset and 0.2852 for the enron dataset) without any need for hyperparameter tuning.

Table 2 shows that, CascadeML is very competitive across different datasets compared to state-of-the-art algorithms while not requiring hyperparameter tuning.

Table 3: Summary of the trained CascadeML network architecture for all datasets.

	Cascade Depth	Scaled hidden nodes
flags	8.30 ± 2.87	0.47 ± 0.15
yeast	3.80 ± 0.79	1.16 ± 0.22
scene	5.10 ± 1.85	0.89 ± 0.28
emotions	5.00 ± 2.58	1.04 ± 0.32
medical	6.80 ± 3.16	0.80 ± 0.20
enron	3.60 ± 0.70	1.16 ± 0.21
birds	7.60 ± 1.58	0.64 ± 0.14
genbase	8.40 ± 3.24	1.04 ± 0.49
cal500	6.40 ± 2.84	1.34 ± 0.60
llog	8.10 ± 3.45	0.67 ± 0.20

CascadeML learns architectures with different number of nodes and activation functions per layer as shown in Table 3 and Figure 3 for every dataset over multiple folds. In Table 3 *Cascade depth* indicates the depth of the cascade network, and *Scaled hidden nodes* indicate the number of total hidden units divided by the number of input units for each dataset. Figure 3a shows the boxplots of the learnnet network depths over folds for each dataset and Figure 3b shows the boxplots for the scaled hidden nodes. Note that although the standard deviations of the performances in Table 2 are small, the trained layer depth and the scaled hidden nodes have high standard deviations. Figure 3c shows a scatterplot of the depths and the scaled hidden nodes values over all datasets and folds. This indicates that the learned networks were either deep with fewer nodes per layer, or shallow but more nodes per layer, therefore having a similar network capacity and hence the F-Score performance over the folds were similar, although the architecture learned were very different.

A network architecture example learned by CascadeML on the yeast dataset is shown in Figure 4. For this specific execution, three cascaded layers were selected with L1 having 220 units and a *tanh* activation, L2 having 64 units and a *linear* activation, and L3 having 26 units and a *linear* activation. yeast

6 Conclusion and Future Work

The work introduces a neural network algorithm, CascadeML, for multi-label classification based on the cascade architecture, which grows the architecture as it trains and takes label associations into account. Except setting some bounds of the hyperparameters, the method omits the requirement of hyperparameter tuning as it automatically determines the architecture, and uses an adaptive first order gradient descent algorithm, iRProp-.

In an evaluation experiment CascadeML was shown to perform competitively to state-of-the-art multi-label classification algorithms, where all the other multi-label algorithms were hyperparameter tuned. CascadeML performed better on an average classifier chains, HOMER with RBF-SVM, BPMLL and MLkNN. RAKEL had the same overall average rank compared to CascadeML, but it did not require the extensive hyperparameter tuning.

Fig. 4: An example network generated by a run of CascadeML on the yeast dataset. The rectangles represents layers, labelled with the number of nodes and the selected activation (act) function the layers. The lines connecting the layers indicate full connection and the text indicates the number of weights involved in the corresponding connection.

CascadeML is the first automatic neural network algorithm with a competitive performance to hyperparameter tuned state-of-the-art multi-label classification methods, although CascadeML’s performance can be improved in the cases where it does perform poorly. A limitation of the BPMLL loss function used in CascadeML is that it cannot scale with increasing number labels [17]. As the comparisons are pairwise, as the number of labels increase the computation becomes slow like BPMLL. Therefore, it would be interesting to investigate alternative loss functions that can still take account of label associations without the need for expensive pairwise comparisons. Also, it would be interesting to examine the patterns in which layers grow during CascadeML so as different mechanisms for adding new layers could be introduced.

References

1. Baluja, S., Fahlman, S.: Reducing network depth in the cascade-correlation learning architecture. Tech. Rep. CMU-CS-94-209, Carnegie Mellon University, Pittsburgh, PA (October 1994)
2. Charte, F., Rivera, A., del Jesus, M.J., Herrera, F.: Concurrence among imbalanced labels and its influence on multilabel resampling algorithms. In: Polycarpou, M., de Carvalho, A.C.P.L.F., Pan, J.S., Woźniak, M., Quintian, H., Corchado, E. (eds.) Hybrid Artificial Intelligence Systems. pp. 110–121. Springer International Publishing, Cham (2014)

3. Chen, Z., Chi, Z., Fu, H., Feng, D.: Multi-instance multi-label image classification: A neural approach. *Neurocomputing* **99**, 298 – 306 (2013)
4. Cheng, W., Hullermeier, E.: Combining instance-based learning and logistic regression for multilabel classification. *Machine Learning* **76**(2-3), 211–225 (2009)
5. Crammer, K., Singer, Y.: A family of additive online algorithms for category ranking. *J. Mach. Learn. Res.* **3**, 1025–1058 (Mar 2003)
6. Fahlman, S.E.: An empirical study of learning speed in back-propagation networks. Tech. rep. (1988)
7. Fahlman, S.E., Lebiere, C.: The cascade-correlation learning architecture. In: Touretzky, D.S. (ed.) *Advances in Neural Information Processing Systems 2*, pp. 524–532. Morgan-Kaufmann (1990)
8. Feurer, M., Klein, A., Eggenberger, K., Springenberg, J., Blum, M., Hutter, F.: Efficient and robust automated machine learning. In: *Advances in Neural Information Processing Systems 28*, pp. 2962–2970 (2015)
9. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press (2016), <http://www.deeplearningbook.org>
10. Grodzicki, R., Mańdziuk, J., Wang, L.: Improved multilabel classification with neural networks. In: Rudolph, G., Jansen, T., Beume, N., Lucas, S., Poloni, C. (eds.) *Parallel Problem Solving from Nature – PPSN X*. pp. 409–416. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
11. Hansen, L.K., Pedersen, M.W.: Controlled growth of cascade correlation nets. In: Marinaro, M., Morasso, P.G. (eds.) *ICANN '94*. pp. 797–800. Springer London, London (1994)
12. Herrera, F., Charte, F., Rivera, A.J., del Jesús, M.J.: *Multilabel Classification - Problem Analysis, Metrics and Techniques*. Springer (2016)
13. Igel, C., Hüsken, M.: Improving the rprop learning algorithm. In: *Proceedings of the second international ICSC symposium on neural computation (NC 2000)*. vol. 2000, pp. 115–121. Citeseer (2000)
14. Kelleher, J.D., Mac Namee, B., D’Arcy, A.: *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies*. The MIT Press (2015)
15. Madjarov, G., Kocev, D., Gjorgjevikj, D., Džeroski, S.: An extensive experimental comparison of methods for multi-label learning. *Pattern Recognition* **45**(9), 3084 – 3104 (2012)
16. Mencia, E.L., Furnkranz, J.: Pairwise learning of multilabel classifications with perceptrons. In: *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. pp. 2899–2906 (June 2008)
17. Nam, J., Kim, J., Loza Mencia, E., Gurevych, I., Fürnkranz, J.: Large-scale multilabel text classification — revisiting neural networks. In: Calders, T., Esposito, F., Hüllermeier, E., Meo, R. (eds.) *Machine Learning and Knowledge Discovery in Databases*. pp. 437–452. Springer Berlin Heidelberg (2014)
18. Nissen, S.: Large scale reinforcement learning using q-sarsa (λ) and cascading neural networks. Unpublished masters thesis, Department of Computer Science, University of Copenhagen, København, Denmark (2007)
19. Pakrashi, A., Greene, D., Mac Namee, B.: Benchmarking multi-label classification algorithms. In: *24th Irish Conference on Artificial Intelligence and Cognitive Science (AICS'16)*, Dublin, Ireland, 20-21 September 2016. CEUR Workshop Proceedings (2016)
20. Phatak, D.S., Koren, I.: Connectivity and performance tradeoffs in the cascade correlation learning architecture. *IEEE Transactions on Neural Networks* **5**(6), 930–935 (Nov 1994)

21. Prechelt, L.: Investigation of the cascor family of learning algorithms. *Neural Networks* **10**(5), 885 – 896 (1997)
22. Read, J., Pérez-Cruz, F.: Deep learning for multi-label classification. *CoRR abs/1502.05988* (2015)
23. Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Machine Learning* **85**(3), 333–359 (2011)
24. Rojas, R.: *Neural Networks: A Systematic Introduction*. Springer-Verlag, Berlin, Heidelberg (1996)
25. de Sá, A.G.C., Freitas, A.A., Pappa, G.L.: Automated selection and configuration of multi-label classification algorithms with grammar-based genetic programming. In: *PPSN* (2018)
26. de Sá, A.G.C., Pappa, G.L., Freitas, A.A.: Towards a method for automatically selecting and configuring multi-label classification algorithms. In: *GECCO* (2017)
27. Spyromitros, E., Tsoumakas, G., Vlahavas, I.: An empirical study of lazy multilabel classification algorithms. In: *Proc. 5th Hellenic Conference on Artificial Intelligence (SETN 2008)* (2008)
28. Tsoumakas, G., Katakis, I., Vlahavas, I.: Effective and efficient multilabel classification in domains with large number of labels. In: *Proc. ECML/PKDD 2008 Workshop on Mining Multidimensional Data (MMD’08)*. vol. 21, pp. 53–59. sn (2008)
29. Tsoumakas, G., Spyromitros-Xioufis, E., Vilcek, J., Vlahavas, I.: Mulan: A java library for multi-label learning. *Journal of Machine Learning Research* **12**, 2411–2414 (2011)
30. Tsoumakas, G., Vlahavas, I.P.: Random k -labelsets: An ensemble method for multilabel classification. In: *ECML* (2007)
31. Waugh, S., Adams, A.: Connection strategies in cascade-correlation. In: *The Fifth Australian Conference on Neural Networks*. pp. 1–4 (1994)
32. Wei, Y., Xia, W., Huang, J., Ni, B., Dong, J., Zhao, Y., Yan, S.: CNN: single-label to multi-label. *CoRR abs/1406.5726* (2014)
33. Wever, M., Mohr, F., Hüllermeier, E.: Automated multi-label classification based on ML-Plan. *CoRR abs/1811.04060* (2018)
34. Yu, Q., Wang, J., Zhang, S., Gong, Y., Zhao, J.: Combining local and global hypotheses in deep neural network for multi-label image classification. *Neurocomputing* **235**, 38 – 45 (2017)
35. Zhang, M.L., Zhou, Z.H.: ML-kNN: A lazy learning approach to multi-label learning. *Pattern Recognition* **40**, 2038–2048 (2007)
36. Zhang, M.L.: ML-RBF: Rbf neural networks for multi-label learning. *Neural Processing Letters* **29**(2), 61–74 (Apr 2009)
37. Zhang, M.L., Zhou, Z.H.: Multilabel neural networks with applications to functional genomics and text categorization. *IEEE Trans. on Knowl. and Data Eng.* **18**(10), 1338–1351 (Oct 2006)
38. Zhang, M.L., Zhou, Z.H.: A review on multi-label learning algorithms. *IEEE transactions on knowledge and data engineering* **26**(8), 1819–1837 (2014)
39. Zhu, J., Liao, S., Lei, Z., Li, S.Z.: Multi-label convolutional neural network based pedestrian attribute classification. *Image and Vision Computing* **58**, 224 – 229 (2017)
40. Zhuang, N., Yan, Y., Chen, S., Wang, H., Shen, C.: Multi-label learning based deep transfer neural network for facial attribute classification. *Pattern Recognition* **80**, 225 – 240 (2018)