ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1.
Available from http://www.i6doc.com/en/.

ELM Preference Learning for Physiological Data

Davide Bacciu[1], Michele Colombo[1,2], Davide Morelli[1,2,3] and David Plans[2,3]

1- Dipartimento di Informatica - Università di Pisa - Italy
2- BioBeats Group Ltd - London - UK
3- Center for Digital Economy - University of Surrey - UK

Abstract. The work confronts two approaches to realize preference learning using Extreme Learning Machine networks, relaying on limited and subject-dependant information concerning pairwise relations between data samples. We describe an application within the context of assessing the effect of breathing exercises on heart-rate variability, using a dataset of over 19K exercising sessions. Results highlight the importance of using weight sharing architectures to learn smooth and generalizable complete orders induced by the preference relation.

1 Introduction

Heart Rate Variability (HRV) has been shown to be related to several physiological aspects including the activity of the Autonomic Nervous System (ANS) [1] which, among the others, influences the relaxation level of humans. Deep diaphragmatic breathing is known to affect HRV [2]. However, learning a direct correlation between HRV activity and ANS activation is a difficult task given the subjective nature of HRV changes. As a result, typical supervised learning approaches fail to capture such a correlation [3]. An alternative way to tackle the problem addresses it as a preference learning task, where the learning model is provided with partial information according to some subjective preferential ordering between samples, e.g. the fact that two HRV samples from the same subject are taken under different ANS conditions. Such partial, pairwise information is then used to build a total ordering of the samples from all subjects according to the unknown total ordering function, e.g. the relaxation level. In this work, we discuss how the problem of learning preferential rankings can be effectively addressed by Extreme Learning Machines (ELMs) [4] by introducing a specialized preference learning objective function coupled with a weight sharing architecture. We provide an experimental assessment of different ELM architectures for preference learning and we show how certain architectural aspects, such as weight sharing, have desirable effects on the realization of the total ordering, despite yielding to slightly less accurate performance on predicting pairwise partial orderings. The experimental assessment exploits heart-rate samples acquired by Biobeats'[1] app *Hear and Now*. Particular care is taken to show generalization of the results to completely unseen subjects, also by relaying on additional external datasets collected under different conditions and throughout different wearable devices.

[1] `http://www.biobeats.com`

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1. Available from http://www.i6doc.com/en/.

2 Extreme learning machines for preference learning

ELMs are a family of neural networks for vectorial data based on the concept of building an high dimensional basis expansion of input information through one (or more) hidden layer of randomly initialized and non-plastic neurons, followed by an output layer of plastic neurons which are adapted to provide predictions in accordance to the task at hand. The random untrained nature of the hidden neurons makes training of the ELM computationally efficient, e.g. when dealing with Big Data, while they retain universal approximation capabilities [4]. The simplest form of ELM network comprises an hidden layer of sigmoid-like neurons which receive in input the data samples and that are connected to an output layer of linear neurons whose number depends on the task at hand. Each output neuron realizes a weighted combination of the hidden neuron outputs to compute the network predictions. In such a model, input weights are typically initialized by random Gaussian noise with zero mean and given variance and left unchanged, while output weights can be trained by standard techniques such as Moore-Penrose pseudo-inverse or its regularized variants.

ELM networks have found application to several classification and regression problems, including a recent extension in the context of learning to rank [5]. Here ELMs are used to learn an ordering between samples $\mathbf{x}_j^i$ representing the relevance of a document j to a query i using supervised target information y_j^i concerning document-query relevance. This setting is similar to what we want to achieve in this paper, except that in our preference learning setting we do not rely on the availability of a relevance value y_i^j to be used as target for the ELM network output. Instead, we consider a set of M input samples $\mathcal{D} = \mathbf{x}_1, \ldots, \mathbf{x}_M$ that is provided together with partial supervised information concerning pairwise ordering between them. For instance, the dataset might provide supervised information stating that $\mathbf{x}_1 \leq_p \mathbf{x}_2$ and $\mathbf{x}_1 \leq_p \mathbf{x}_3$, where $\leq_p$ is the preference relation we want to learn, but nothing is stated concerning the relationship between $\mathbf{x}_2$ and $\mathbf{x}_3$. Training a model using only information on pairwise preferences without having it fit to specific relevance/preference values (e.g. y_j^i in the ranking application) is an advantage whenever it does not exist a uniform or consolidated preference function f_p across the samples. This is the case when the preference function is subjective to the samples such as when a sample represents an individual-related measurement or a self-assessment [3]. In such cases, we cannot expect different individuals to self assess using the same reference scoring functions, hence the necessity of avoiding to use the value of the preference function as a target for the learning model. On the other hand, the preference value can be used to assess the pairwise ordering $\leq_p$ between samples taken from the same individual and which, thus, share the same scoring system. In the following, we discuss two architectural variants of ELM to realize such a preference learning task.

A first, straightforward, way to approach such a problem is to cast it as a classification task where an ELM network receives in input two samples, $\mathbf{x}_1$ and $\mathbf{x}_2$, and it is trained to predict a binary value determining which sample is to be

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1. Available from http://www.i6doc.com/en/.

preferred over the other. More formally, the (scalar) response of the network to a joint input $\mathbf{x}^i = [\mathbf{x}_1, \mathbf{x}_2]$ is interpreted as follows

$$y^i_{1,2} = \begin{cases} +1\text{:} & \mathbf{x}_1 \text{ ranks before or equal to } \mathbf{x}_2 \\ -1\text{:} & \mathbf{x}_2 \text{ ranks before than } \mathbf{x}_1 \end{cases}. \tag{1}$$

Such pairwise preference problem is a standard binary classification task on the joint input $\mathbf{x}^i$ with binary target y^i to be interpreted as in (1). A second approach to the task stems from the consideration that the formulation above neglects one key aspect of the problem, that is symmetry. Supposing that $\mathbf{x}_1$ is to be preferred to $\mathbf{x}_2$ this means that it can be presented to the network as the joint sample $\mathbf{x}^i = [\mathbf{x}_1, \mathbf{x}_2]$ with target $+1$ as well as the joint sample $\mathbf{x}^{i'} = [\mathbf{x}_2, \mathbf{x}_1]$ with target -1. This aspect suggests that we can put in place a weight sharing approach where a single ELM network computes its output $\hat{f}(\mathbf{x}_j)$ for each sample $\mathbf{x}_j \in \mathcal{D}$, flipping the sign of $\hat{f}(\mathbf{x}_j)$ depending on whether $\mathbf{x}_j$ is used as first or second term of the pairwise comparison. This can be interpreted as an extended ELM network with two modules sharing identical $\mathbf{W}_{in}$ weights and sign-complemented $\mathbf{W}_{out}$ weights. Training of a such a network can be achieved by minimization of the following preference learning error [3] with respect to $\mathbf{W}_{out}$:

$$E(\mathcal{D}) = \sum_{p,j} \frac{1}{2} \| 2 - (\hat{f}(\mathbf{x}_p) - \hat{f}(\mathbf{x}_j)) \|^2, \tag{2}$$

where $\mathbf{x}_p$ and $\mathbf{x}_j$ are sample pairs in $\mathcal{D}$ where $\mathbf{x}_p$ is to be preferred over $\mathbf{x}_j$. Note how the error function in (2) does not require a target output for the network, whereas the sole preference relation information is used to separate the network outputs for those sample pairs for which a preference relation exists.

3 An application of preference learning to data from mobile heart-rate sensors

The two approaches described in the previous section have been applied to learning pairwise preferential rankings between heart-rate samples to learn to assess the effect of breathing exercises on target users. Biobeats released an app for iOS devices, called Hear and Now[2], that guides users through slow diaphragmatic breathing exercises. In doing so, the app also collects the heart rate variability for 40 seconds before the breathing exercise, and for 40 seconds after the breathing exercise, to measure the effect of the breathing exercise. By these means, we have composed a dataset of 19567 sessions providing pairwise measurements of 10 HRV features before and after each breathing session. These are popular HRV features including, e.g., the standard deviation of the time between heartbeats (SDNN) [6] and the root mean square of successive differences of heartbeats (RMSSD) [1]. Such data has been used as a benchmark to assess the two preference learning approaches in terms of predictive performance, i.e.

[2] https://itunes.apple.com/us/app/hear-and-now/id977650202

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1. Available from http://www.i6doc.com/en/.

the ability of telling pre-breathing samples from post ones, as well as in terms of quality of the global order that can be constructed by application of the learned pairwise comparisons. A sample of 15% of the original users has been randomly selected with all associated sessions to be used as an external test population to assess the generalization to completely unknown users. By this means, we have composed an hold out test set of 2637 sessions. The remainder of the data has been split into 12079 and 4144 training and validation sessions, respectively.

The experimental analysis compares the two architectural variants in Section 2 referred to as *ELM-WS* and *ELM-BIN*, respectively, for the model with and without weight sharing. Both networks have a single hidden layer and hyperbolic tangent activation both for the hidden and the output neurons, while output weights are trained by stochastic gradient descent with momentum (given the nonlinearity of the output neurons) applied to a mean-squared error loss. Model selection has been performed to select the best hyperparameterization on the validation set, considering hyperparameters such as hidden layer size ($[5000, 7500, 10000]$), learning rate ($[0.01, 1e^{-3}, 1e^{-4}, 1e^{-5}]$), momentum ($[0, 0.3, 0.6]$), error regularization ($[0, 1e^{-2}, 1e^{-3}, 1e^{-4}]$). Three networks have been trained for each configuration and their results averaged. Input data has been normalized by z-scoring.

Table 1 shows the accuracy of the different models in terms of correctly predicted pairwise preferences. Note that, in our preference relation, we assume that the HRV measurement before the breathing exercise (pre) should precede the corresponding HRV after the exercise (post). The intended aim is to exploit preference learning to try to capture an unknown global ordering between HRV measurements based on subject-dependent information on episodic relaxation induced by the breathing exercise. To provide a baseline performance for the ELM, Table 1 also shows the results for an MLP model trained for binary classification (like ELM-BIN) with hidden layer size in $[10, 75]$. Results highlight that the models are able to generalize the preference relation learned on users in training to unknown subjects in test (cf. the $\approx 4\%$ difference between validation and test set). In particular, the ELM-BIN model seems to have the highest performance in separating pre from post breathing HRV measurements, while the effect of the weight sharing architecture does not seem beneficial this stage. However, being able to detect the correct pairwise preference relation between two HRV measurements from the same session and the same user is not the ultimate goal of our application. Rather, we wish to be able to use the learned pairwise relation to determine a complete ordering of the HRV samples according to the, unknown, joint preference relation across all the users.

	ELM-WS		ELM-BIN		MLP	
	val	*tst*	*val*	*tst*	*val*	*tst*
Acc. (%)	77.1 (0.7)	72.8 (0.3)	77.8 (2.0)	74.0 (0.8)	77.0 (0.6)	73.8 (0.6)

Table 1: Accuracy in predicting the pairwise preference on the *Biobeats* data: results refer to the best model selected in validation; values are averaged on the 3 model initializations and standard deviation is in brackets.

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1.
Available from http://www.i6doc.com/en/.

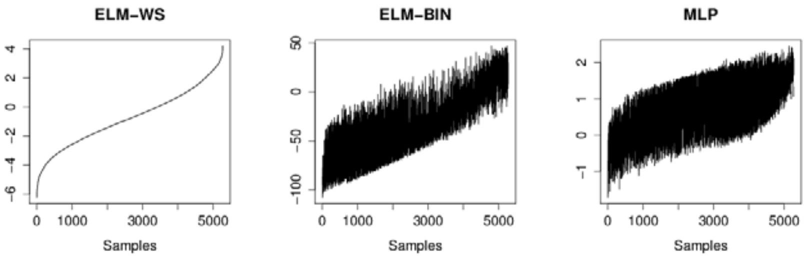

Fig. 1: Concordance plot of the orderings generated by pairwise comparisons between random pivotal HRV measurements and the rest of the samples in the test-set for ELM-WS, ELM-BIN and MLP (from left to right).

In this particular application, this accounts to be able to rank the HRV samples in order from the one corresponding to maximum breathing effect to the one corresponding to minimum breathing effect. In this respect, a *good* learned relation is the one that provides an ordering where the pairwise comparisons are in agreement with each other. Figure 1 shows an example of such a quality measure for the 3 models in Table 1. Each plot is generated by pairwise comparisons of pre or post HRV samples with two randomly selected pivotal measurements (i.e. maintaining one term of the comparison fixed to the pivot and varying the other across all dataset). Samples are ordered on the x-axis accordingly to a pivot measurement, while y-axis values are taken accordingly to the other one. The smoother the line, the more concordant the two orderings generated by the two pivots. The results in Fig. 1 show that ELM-WS yields to orderings that are extremely concordant among pivots thanks to the symmetry introduced by the weight sharing architecture, while both ELM-BIN and the MLP yield to poor quality total preference orderings, despite being sightly more precise than ELM-WS in deciding the outcome of a pairwise confrontation (see Table 1).

The significant aspect of the weight sharing architecture seems to be the ability to generalize subjective pairwise relationships to a common total relationship shared by multiple subjects who have no sessions in the training set. To further assess this aspect and to show a potential application of the approach, we have applied the ELM-WS trained on the Biobeats dataset to an external test set collected by Biobeats in collaboration with a corporate partner providing a wrist worn band measuring heart rate 24/7. Continuous heart rate was collected from approximately 500 users over 2 months, while users were also provided with a customized version of Hear and Now guiding them in slow breathing exercises. For every user, HRV features have been computed for times up to 120 minutes before and after each breathing session: Fig. 2 shows the average output predicted by ELM-WS before and after the breathing exercise: here ELM-WS is shown to detect the effects of the breathing exercise, encoded by output values tending more towards -1. Such result provides an indication of how the ELM-WS approach can generalize to novel devices (i.e. wristband hearts sensor in place of mobile camera) and to completely unknown subject cohorts.

ESANN 2017 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium), 26-28 April 2017, i6doc.com publ., ISBN 978-287587039-1.
Available from http://www.i6doc.com/en/.

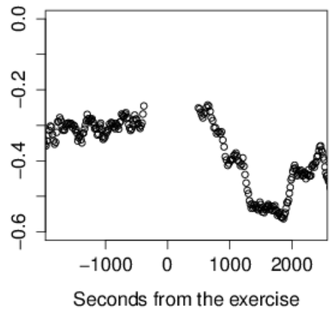

Fig. 2: ELM-WS output (on the y-axis) for sec-onds preceding and follow-ing a breathing exercise which is performed in correspondence of the area without markers around the zero of the x-axis. The plot aligns and averages multiple breathing exercise episodes from 500 users from the wristband pilot.

4 Conclusion

We have assessed the realization of a preference learning task on HRV features exploiting two architectural variants of ELM networks. Our analysis suggests that using weight sharing techniques to exploit symmetries in the task induces a smoother total preference relation, learned from potentially very different personal pairwise confrontations. Such performance generalizes very well also to external data taken under different experimental conditions, being able to track and predict the execution of breathing exercises across the day. Future work concerns exploring more powerful ELM architectures such as bidirectional [7] and fully complex ELMs. Ongoing work is also exploring preference learning in the context of another class of random networks from the reservoir computing paradigm [8] which allows to process the HRV information in its raw sequential form, without the need of extracting engineered features in vectorial form.

References

[1] E. S. Mezzacappa, R. M. Kelsey, E. S. Katkin, and R. P. Sloan. Vagal rebound and recovery from psychological stress. *Psychosomatic Medicine*, 63(4):650–657, 2001.

[2] G. K. Pal, S. Velkumary, and others. Effect of short-term practice of breathing exercises on autonomic functions in normal human volunteers. *Indian Journal of Medical Research*, 120(2):115, 2004.

[3] H. P. Martínez, G. N. Yannakakis, and J. Hallam. Don't classify ratings of affect; rank them! *IEEE Transactions on Affective Computing*, 5(3):314–326, July 2014.

[4] G.B. Huang, Q.Y. Zhu, and C. K. Siew. Extreme learning machine: Theory and applications. *Neurocomputing*, 70(1–3):489–501, 2006.

[5] W. Zong and G. B. Huang. Learning to rank with extreme learning machine. *Neural Processing Letters*, 39(2):155–166, 2014.

[6] J. P. A. Delaney and D. A. Brodie. Effects of short-term psychological stress on the time and frequency domains of heart-rate variability. *Perceptual and Motor Skills*, 91(2):515–524, 2000.

[7] Y. Yang, Y. Wang, and X. Yuan. Bidirectional extreme learning machine for regression problem and its learning effectiveness. *IEEE Transactions on Neural Networks and Learn Systems*, 23(9):1498–1505, 2012.

[8] D. Bacciu, P. Barsocchi, S. Chessa, C. Gallicchio, and A. Micheli. An experimental characterization of reservoir computing in ambient assisted living applications. *Neural Computing and Applications*, 24(6):1451–1464, 2014.