

Neural Networks in High Energy Physics

Roberto Odorico

University of Bologna, Department of Physics
and
Istituto Nazionale di Fisica Nucleare, Sezione di Bologna
Via Imerio 46, 40126 Bologna, Italy
(e-mail: odorico@bologna.infn.it)

ABSTRACT

A short survey of the use of Neural Networks and statistical discriminants in High Energy Physics is presented. The focus is on classification problems, e.g. involving events or jets. After illustrating the various neural and statistical classifiers currently used, some assessment of their comparative performance for top and bottom jets is made.

Neural Networks (NN) are receiving increasing attention for t-quark [1-3] and b-quark [4] tagging with $t, b \rightarrow$ jets. Also of interest are other applications like e/π discrimination for improved lepton tagging. Especially used are NN based on Learning Vector Quantization [5] to which Training Count can be profitably added (LVQTC) [3,6] and Back-Propagation (BP) [7-9]. One should not forget more conventional statistical methods like the Fisher linear discriminant [10,11] and the Gaussian classifier [11].

To get some orientation, let me start with the simplest classifier: Fisher's linear discriminant [10,11]. As for all classifiers, one must first encode the event (or whatever objects one is classifying, e.g. jets) into a number of feature variables s_1, s_2, \dots, s_n , which can be arranged into a vector, the pattern vector s . Thus, each event corresponds to a point s in the multidimensional pattern space. Let us consider a schematic 2-dimensional example: Fig. 1. When projecting the events onto the s_1 axis, one gets two overlapping distributions in s_1 . When trying to discriminate the two classes by a cut in s_1 one gets penalized in efficiency and purity. The same holds true when projecting onto the s_2 axis. But if one projects onto the Fisher axis shown in the figure, the two distributions get separated, and thus discrimination with 100% purity and efficiency is achieved. That is a dream situation, of course, but in a real case one can reduce the overlap between the two distributions to some minimal value in this way. The Fisher variable F , associated with the Fisher axis, is a linear combination of the feature variables s_1, s_2, \dots, s_n , which can be determined by simple mathematics involving the the correlation (or covariance) matrices for the two distributions. The physical ingredient one is exploiting is the **correlation** between the variables within each class distribution. In the example considered, for each s_2 bin s_1 ranges over two distinct intervals for the two distributions. As the s_2 bin moves, the populated s_1 intervals change, which means that s_1 is correlated to s_2 within each distribution. The s_1 - s_2 correlation turns out to be different within each distribution, and that is what is exploited in the discrimination.

The Fisher classifier cannot be outperformed if the two distributions are gaussian and have the same correlation matrix, i.e. if they have the same shapes and just distinct centroids. In this case, its purity and efficiency of classification are only limited by the amount of overlap between the two distributions. If the two distributions have different correlation matrices, one can introduce a more general Gaussian classifier, that cannot be outperformed if the two distributions are gaussian [11]. For that, one considers the probability density functions corresponding to the gaussian approximations for the two distributions:

$$p(s) = A \exp\left\{-\frac{1}{2}(s-\langle s \rangle)M^{-1}(s-\langle s \rangle)\right\}$$

where M is the correlation matrix for the distribution and $\langle s \rangle$ is its centroid. The Gaussian classifier for a pattern s is given by the variable $G = \ln(p_A(s)/p_B(s))$. If G is positive, s is classified as A, otherwise as B. The absolute magnitude of G gauges the **reliability** of the

classification. If the two correlation matrices are equal, G identifies with the Fisher variable F apart from a constant term: $G = F + \ln(A_A/A_B)$.

If the distributions are not gaussian, in general the Gaussian classifier does not yield the best purity and efficiency which are in principle obtainable given the overlap between the two distributions (i.e. it no longer reaches the Bayesian limit). That is where Neural Networks come to help.

As an illustration, let us consider the simple 2-dimensional example of Fig. 2, with the two distributions being uniform within the regions they cover. The outer distribution (B) is quite far from gaussian. The Gaussian classifier gives a bad performance in this case.

Let us see how a NN like LVQTC [5,3,6] handles the problem. With this NN architecture neurons can be associated with vectors, or points, in pattern space. Their positions are fixed by a training procedure in which a sequence of patterns of known class $s(t)$, $t = 1, 2, 3, \dots$ is presented. For each pattern $s(t)$, one corrects the position of the neuron closest to it, m_c , by moving it closer to the pattern if the two belong to the same class, or moving it away from it if they belong to different classes:

$$\begin{aligned} m_c(t+1) &= m_c(t) + \alpha_+(t) [s(t) - m_c(t)] && \text{if } m_c \text{ and } s \text{ belong to the same class} \\ m_c(t+1) &= m_c(t) - \alpha_-(t) [s(t) - m_c(t)] && \text{if } m_c \text{ and } s \text{ belong to different classes} \end{aligned}$$

$\alpha_+(t)$ and $\alpha_-(t)$ are positive learning parameters decreasing with t . During training the number of times each neuron is corrected by patterns of the various classes is counted. From that the neuron purity can be calculated, i.e. the fraction of times the neuron is corrected by patterns of its own class. The classification of a pattern s of unknown class is made by simply assigning the pattern to the class of its closest neuron. The purity of the classification can be estimated by the purity of the neuron providing the classification. For the problem of Fig. 2, 100% purity and efficiency of classification can be achieved with 1 neuron of class A and 16 neurons of class B. Their positions in Fig. 2 are those resulting from training.

In a BP net [7] the relevant neurons are not associated with points but rather with hyperplanes in pattern space. For the 2-dimensional problem of Fig. 2, the hyperplanes become just straight lines. In such a net, neurons are arranged in successive layers, the excitations of neurons in a layer being determined by the excitations of neurons in the previous layer, Fig. 3 (hence the name of Multi-Layered Perceptron more correctly used for such a net, the term Back-Propagation referring more properly to the type of training algorithm used). The excitations of neurons in the input layer ($L=0$), one for each pattern component, are directly given by the values of the corresponding pattern components. The excitation $x_i(L=1)$ of a neuron in the next hidden layer is determined by feeding the value of the linear expression $a_i = \sum_k \omega_{ik} s_k + \theta_i$ into a saturating transfer function: $x_i(L=1) = g(a_i)$, e.g. $g(a_i) = \tanh(a_i)$. a_i can be visualized as the distance of the pattern from a straight line (a

hyperplane in the general case), whose orientation and position are determined by the "weight" parameters ω_{ik} and the "bias" term θ_i associated with the neuron. Several hidden layers may be included, but one has been found to be enough in most HEP applications. The last, output, layer may consist of several neurons, but one is enough if the classes are just 2. Its excitation is determined by iterating the procedure used to calculate the excitations of the hidden neurons. The value of the net is due to the existence of a training algorithm (BP) which starting from the discrepancies of the output excitations with respect to the desired (target) results for each training pattern (of known class) corrects the net parameters (weights and bias terms) so as to achieve minimum output discrepancies at the end of training. There is no guarantee, though, that the minimum obtained is an "absolute" minimum. For our example, 100% purity and efficiency of classification are achieved on an independent test set of patterns by using 3 hidden neurons, represented by the 3 straight lines in Fig. 2, and an output excitation given by $x(L=2) = \tanh(1.5 \sum_i x_i(L=1) + 4.5)$ (the distance metric from the hidden neurons lines is 4.5 times the Euclidean metric). A pattern is classified as A if $x(L=2) < 0.5$, otherwise as B.

The example helps to illustrate some important differences in the usage of LVQTC and BP nets.

BP requires a relatively limited number of parameters and thus the training statistics can be kept small. The cpu time required for training is typically long, since when correcting positions of hyperplanes describing hidden neurons to better accommodate patterns in a given region of pattern space, far away patterns can easily get penalized. Optimization must thus be handled globally, going through the whole training set. The purity/efficiency trade-off in classification can be controlled by cuts on the output excitation.

LVQTC requires comparatively many more parameters, and therefore the training statistics must be much larger than in BP. The hoped for reward is a higher degree of purity in classification. The cpu training time is typically short, since in order to better accommodate patterns in a region one must correct only neurons in that region, without affecting classification performance for far away patterns. The purity/efficiency trade-off can be controlled by cuts in the neuron purity.

Dedicated hardware implementations of BP and LVQTC (e.g. for triggers) also have distinct requirements. BP needs vectorization and a realization of the non-linear transfer function. LVQTC can profit of massive parallelism of vector processors, the winner-takes-all step being implemented by a few cycles of a neural net with lateral inhibitions.

A typical field of application of statistical and NN classifiers is provided by the discrimination of top and bottom jets, originated by t and b quarks decaying into anything (i.e. without a rate-reducing lepton tagging). In [2] it has been shown that for $t\bar{t}$ events produced at the Fermilab Tevatron collider one can get a signal/background ≈ 1.5 with a residual $\sigma(t\bar{t}) \approx 2$ pb, for $m_t = 100$ GeV, by using Fisher's discrimination after some

preliminary cuts. The utilization of LVQTC [3] or of BP does not improve on this result. During the last year a substantial number of contributions have appeared on the utilization of NN's for discriminating b jets at LEP [4]. For the sake of illustrating the comparative performances of the various classifiers in this case, we have applied them to jets in e^+e^- 2-jet events at LEP, simulated by COJETS [12]. The 25 jet feature variables of Bellantoni et al. [4] have been used. The purity versus efficiency results are shown in Fig. 4. In this particular case, it appears that the LVQTC, BP and Gaussian classifiers yield comparable results, with LVQTC having an edge if one wants to have a high purity sample, at the price of a reduced efficiency.

References

- 1) R. Odorico, Phys. Lett. 120B (1983) 219; G. Ballocci and R. Odorico, Nucl. Phys. B229 (1983) 1
- 2) A. Cherubini and R. Odorico, Z. Physik C 47 (1990) 547
- 3) A. Cherubini and R. Odorico, Z. Physik C 53 (1992) 139; A. Cherubini and R. Odorico, Proc. of the Workshop on "Neural Networks: from Biology to High Energy Physics", June 5-14, 1991, Marciana Marina, Isola d'Elba (Italy), ETS Editrice, Pisa (1991), p. 515
- 4) L. Lönnblad et al., Nucl. Phys. B349 (1991) 675; C. Bortolotto et al., Nucl. Instr. Meth. A306 (1991) 459, Udine preprint 91/04/AA (1991); L. Bellantoni et al. (Wisconsin), CERN-PPE/91-80 (1991); J. Proriot et al. (Clermont-Ferrand), Proc. Elba NN Workshop (1991), p. 419; M. Los and N. DeGroot (NIKHEF-H), Proc. Elba NN Workshop (1991), p. 459, CERN NN Workshop (December 1991); J. Jousset et al. (Clermont-Ferrand), L'Agelonde Workshop (January 1992); P. Branchini et al. (INFN Sanità, Rome), L'Agelonde Workshop (January 1992), preprint INFN-ISS 92-1 (1992); F. Block (CERN), CERN NN Workshop (December 1991)
- 5) T. Kohonen, "Self Organization and Associative Memory", 2nd ed., Springer, Berlin (1988)
- 6) A. Cherubini and R. Odorico, LVQNET vers. 1.10, Bologna preprint DFUB 91/13 (1991)
- 7) D.E. Rumelhart, G.E. Hinton and R.J. Williams, "Learning Internal Representations by Error Propagation", in D.E. Rumelhart and J.L. McClelland (Eds.), "Parallel Distributed Processing: Explorations in the Microstructure of Cognition" (Vol. 1), MIT Press (1986), p. 318
- 8) L. Lönnblad et al., JETNET vers. 2.00, Lund preprint LU TP 91-18 (1991)
- 9) J. Proriot, MPL vers. 1.00. Clermont-Ferrand preprint (1991)
- 10) R.A. Fisher, Annals Eugenics 7 (1936) 179
- 11) M. Kendall, A. Stuart and J.K. Ord, "The Advanced Theory of Statistics", Vol. 3, 4th ed., C. Griffin & Co. Ltd., London
- 12) R. Odorico, Comp. Phys. Comm. 32 (1984) 173; COJETS version 6.23, University of Bologna preprint DFUB 91/13 (1991)

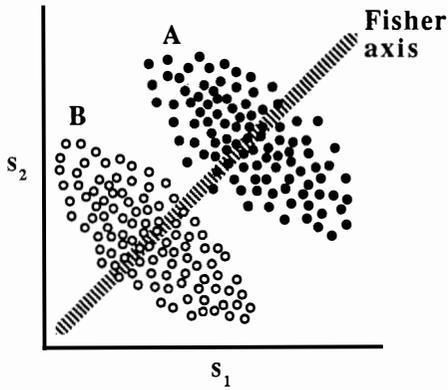


Fig. 1 - Classification example illustrating how Fisher's linear discrimination works.

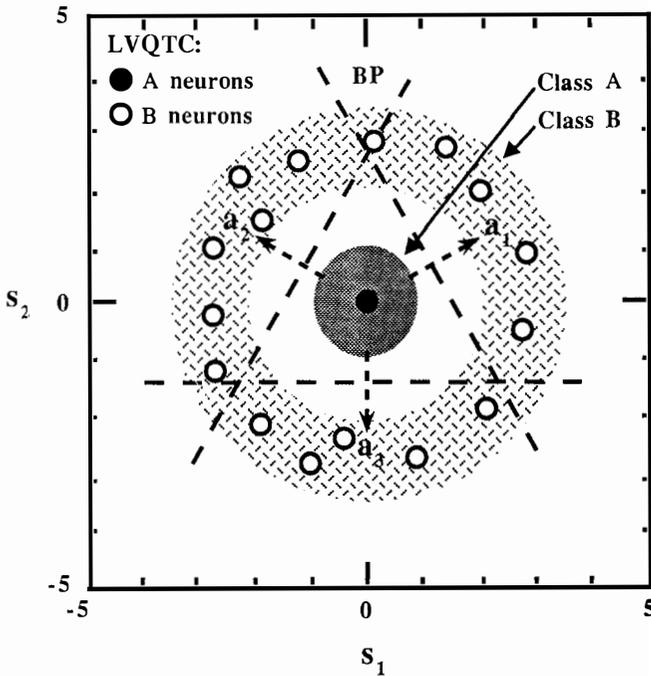


Fig. 2 - Classification example illustrating the way LVQTC and BP neural nets work. Circles represent LVQTC neurons. Straight lines represent BP hidden neurons.

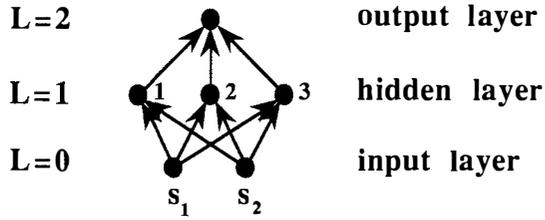


Fig. 3 - Architecture of the BP net used for the example of Fig. 2.

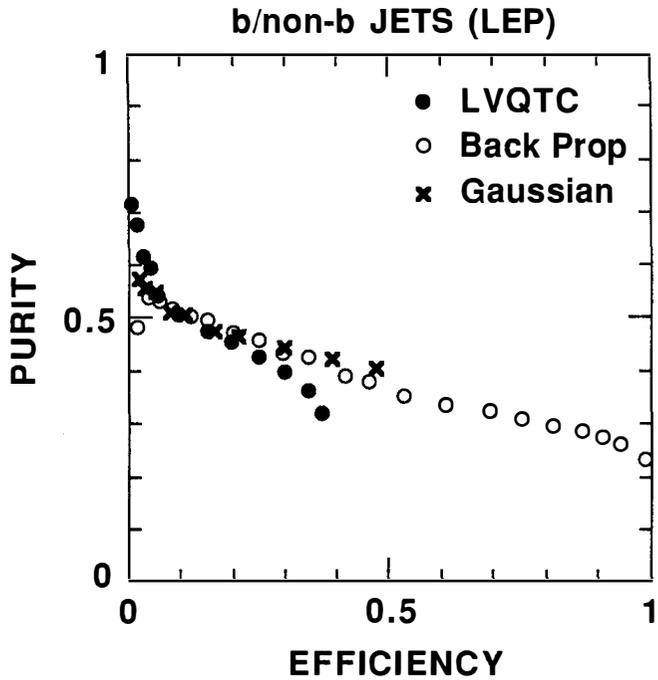


Fig. 4 - Purity vs efficiency for b/non-b jets at LEP according to LVQTC, BP and Gaussian jet classification.