



CONTRIBUTED ARTICLE

Characteristic Functions and Process Identification by Neural Networks

JOAQUIM A. DENTE¹ AND RUI VILELA MENDES^{2,3}

¹Departamento de Engenharia Electrotécnica e de Computadores, Instituto Superior Técnico, , ²Grupo de Física–Matemática, Universidade de Lisboa, Portugal, and ³Centre de Physique Théorique, CNRS, France

(Received 12 October 1995; accepted 12 March 1997)

Abstract—Principal component analysis (PCA) algorithms use neural networks to extract the eigenvectors of the correlation matrix from the data. However, if the process is non-Gaussian, PCA algorithms or their higher order generalisations provide only incomplete or misleading information on the statistical properties of the data. To handle such situations we propose neural network algorithms, with an hybrid (supervised and unsupervised) learning scheme, which constructs the characteristic function of the probability distribution and the transition functions of the stochastic process. Illustrative examples are presented, which include Cauchy and Lévy-type processes. © 1997 Elsevier Science Ltd.

Keywords—Characteristic functions, Neural networks, Transition functions.

1. INTRODUCTION

Let x_i denote the output of node i in a neural network. Hebbian learning (Hebb, 1949) is a type of unsupervised learning where the neural network connection strengths W_{ij} are reinforced whenever the products $x_i x_j$ are large. If Q is the correlation matrix

$$Q_{ij} = \langle x_i x_j \rangle \quad (1)$$

and the Hebbian learning law is local, all the lines of the connection matrix W_{ij} will converge to the eigenvector of Q with the largest eigenvalue. To obtain other eigenvector directions requires non-local laws (Sanger, 1989; Oja, 1989, 1992; Dente & Vilela Mendes, 1996). These principal component analysis (PCA) algorithms find the characteristic directions of the correlation matrix Q . If the data has zero mean ($\langle x_i \rangle = 0$) they are the orthogonal directions along which the data has maximum variance. If the data is Gaussian in each channel, it is distributed as a hyperellipsoid and the correlation matrix Q already contains all information about the statistical properties. This is because higher order

moments of the data may be obtained from the second order moments. However, if the data is non-Gaussian, the PCA analysis is not complete and higher order correlations are needed to characterise the statistical properties. This led some authors (Softky & Kammen, 1991; Taylor & Coombes, 1993) to propose networks with higher order neurons to obtain the higher order statistical correlations of the data. An higher order neuron is one that is capable of accepting, in each of its input lines, data from two or more channels at once. There is then a set of adjustable strengths $W_{ij_1}, W_{ij_2}, \dots, W_{ij_1 \dots j_n}$, n being the order of the neuron. Networks with higher order neurons have interesting applications, for example in fitting data to a high dimensional hypersurface. However there is a basic weakness in the characterisation of the statistical properties of non-Gaussian data by higher order moments. Existence of the moments of a distribution function depends on the behaviour of this function at infinity and it frequently happens that a distribution has moments up to a certain order, but no higher ones. A well-behaved probability distribution might even have no moments of order higher than one (the mean). In addition a sequence of moments does not necessarily determine a probability distribution function uniquely (Lukacs, 1970). Two different distributions may have the same set of moments. Therefore, for non-Gaussian data, the PCA algorithms or higher order generalisations may lead to misleading results.

Reprint requests should be addressed to Rui Vilela Mendes, Grupo de Física–Matemática, Complexo II, Universidade de Lisboa, Av. Gama Pinto 2, 1699 Lisboa Codex, Portugal, and, Centre de Physique Théorique, CNRS, Luminy, Case 907, F 13288 Marseille Cedex 9, France. Tel.: 351 1 7950790; Fax: 351 1 7954288; e-mail: vilela@alf4.cii.fc.ul.pt




FIGURE 1. A two-dimensional test signal.

As an example consider the two-dimensional signal shown in Figure 1. Figure 2 shows the evolution of the connection strengths W_{11} and W_{12} when this signal is passed through a typical PCA algorithm. Large oscillations appear and finally the algorithm overflows. Smaller learning rates do not introduce qualitative modifications in this evolution. The values may at times appear to stabilise, but large spikes do occur. The reason for this behaviour is that the seemingly harmless data in Figure 1 is generated by a linear combination of a Gaussian with the following distribution

$$p(x) = K(2 + x^2)^{-\frac{1}{2}}$$

which has first moment, but no moments of higher order.

To be concerned with non-Gaussian processes is not a pure academic exercise because, in many applications, adequate tools are needed to analyse such processes. For example, processes without higher order moments, in particular those associated with Lévy statistics, are prominent in complex processes such as relaxation in glassy materials, chaotic phase diffusion in Josephson




FIGURE 2. Evolution of the connection strengths W_{11} and W_{12} in a PCA network for the data in Figure 1.

junctions and turbulent diffusion (Shlesinger et al., 1993; Zumofen & Klafter, 1993, 1994).

Moments of an arbitrary probability distribution may not exist. However, because every bounded and measurable function is integrable with respect to any distribution, the existence of the characteristic function $f(\alpha)$ is always assured (Lukacs, 1970).

$$f(\alpha) = \int e^{i\alpha \cdot \mathbf{x}} dF(\mathbf{x}) = \langle e^{i\alpha \cdot \mathbf{x}} \rangle \quad (2)$$

where α and \mathbf{x} are N -dimensional vectors, \mathbf{x} is the data vector and $F(\mathbf{x})$ its distribution function.

The characteristic function is a compact and complete characterisation of the probability distribution of the signal. If, in addition, one wishes to describe the time correlations of the stochastic process $\mathbf{x}(t)$, the corresponding quantity is the characteristic functional (Hida, 1980)

$$F(\xi) = \int e^{i(\mathbf{x}, \xi)} d\mu(\mathbf{x}) \quad (3)$$

where $\xi(t)$ is a smooth function and the scalar product is

$$(\mathbf{x}, \xi) = \int dt \mathbf{x}(t)\xi(t) \quad (4)$$

where $\mu(\mathbf{x})$ is the probability measure over the sample paths of the process.

In the following we develop an algorithm to compute the characteristic function from the data, by a learning process. The main idea is that in the end of the learning process we should have a neural network which is a representation of the characteristic function. This network is then available to provide all the required information on the probability distribution of the data being analysed. To obtain full information on the stochastic process, a similar algorithm might be used to construct the characteristic functional. However this turns out to be computationally very demanding. Instead we develop a network to learn the transition function and from this the process may be characterised.

2. LEARNING THE CHARACTERISTIC FUNCTION

Suppose we want to learn the characteristic function $f(\alpha)$ (eqn (2)) of a one-dimensional signal $x(t)$ in a domain $\alpha \in [\alpha_0, \alpha_N]$. The α -domain is divided into N intervals by a sequence of values $\alpha_0, \alpha_1, \alpha_2, \dots, \alpha_N$ and a network is constructed with $N + 1$ intermediate layer nodes and an output node (Figure 3).

The learning parameters in the network are the connection strengths W_{0i} and the node parameters θ_i . The existence of the node parameter means that the output of a node in the intermediate layer is $\theta_i \chi_i(\alpha)$, χ_i being a non-linear function. The use of both connection strengths and node parameters in neural networks makes them equivalent to a wide range of other connectionist systems




FIGURE 3. Network to learn the characteristic function of a scalar process.

(Doyme Farmer, 1990) and improves their performance in standard applications (Dente & Vilela Mendes, 1996). The learning laws for the network of Figure 3 are:

$$\theta_i(t+1) = \theta_i(t) + \gamma(\cos \alpha_i x(t) - \theta_i(t)) \quad (5)$$

$$W_{0i}(t+1) = W_{0i}(t) + \eta \times \sum_j \left[\theta_j(t) - \sum_k W_{0k}(t) \chi_k(\alpha_j) \theta_k(t) \right] \times \theta_i(t) \chi_i(\alpha_j).$$

$\gamma, \eta > 0$. The intermediate layer nodes are equipped with a radial basis function

$$\chi_i(\alpha) = \frac{e^{-(\alpha - \alpha_i)^2 / 2\sigma_i^2}}{\sum_{k=0}^N e^{-(\alpha - \alpha_k)^2 / 2\sigma_k^2}} \quad (6)$$

where in general we use $\sigma_i = \sigma$ for all i . The output is a simple additive node.

The learning constant γ should be sufficiently small to insure that the learning time is much larger than the characteristic times of the data $x(t)$. If this condition is satisfied each node parameter θ_i tends to $\langle \cos \alpha_i x \rangle$, the real part of the characteristic function $f(\alpha)$ for $\alpha = \alpha_i$.

The W_{0i} learning law was chosen to minimise the error function

$$f(W) = \frac{1}{2} \sum_j \left(\theta_j - \sum_k W_{0k} \chi_k(\alpha_j) \theta_k \right)^2. \quad (7)$$

One sees that the learning scheme is an hybrid one, in the




FIGURE 4. Network to learn the characteristic function of a two-dimensional signal $\vec{x}(t)$.

sense that the node parameter θ_i learns, in an unsupervised way, (the real part of) the characteristic function $f(\alpha_i)$ and then, by a supervised learning scheme, the W_{0i} 's are adjusted to reproduce the θ_i value in the output whenever the input is α_i . Through the learning law (5) each node parameter θ_i converges to $\langle \cos \alpha_i x \rangle$ and the interpolating nature of the radial basis functions guarantees that, after training, the network will approximate the real part of the characteristic function for any α in the domain $[\alpha_0, \alpha_N]$.

A similar network is constructed for the imaginary part of the characteristic function, where now

$$\theta_i(t+1) = \theta_i(t) + \gamma(\sin \alpha_i x(t) - \theta_i(t)). \quad (8)$$

For higher dimensional data the scheme is similar. The number of required nodes is N^d for a d -dimensional data vector $\mathbf{x}(t)$. For example, for the two-dimensional data of Figure 1 we have used a set of N^2 nodes (Figure 4).

Each node in the square lattice has two inputs for the two components α_1 and α_2 of the vector argument of $f(\vec{\alpha})$. The learning laws are, as before

$$\theta_{(ij)}(t+1) = \theta_{(ij)}(t) + \gamma(\cos(\vec{\alpha}_{(ij)} \cdot \vec{x}(t)) - \theta_{(ij)}(t)) \quad (9)$$

$$W_{0(ij)}(t+1) = W_{0(ij)}(t) + \eta \sum_{(kl)} \left[\theta_{(kl)}(t) - \sum_{(mn)} W_{0(mn)} \chi_{(mn)}(\vec{\alpha}_{(kl)}) \theta_{(mn)}(t) \right] \times \theta_{(ij)}(t) \chi_{(ij)}(\vec{\alpha}_{(kl)}).$$

The pair (ij) denotes the position of the node in the square lattice and the radial basis function is




FIGURE 5. (a) Real part of the characteristic function for the data in Figure 1 (left) and the mesh of θ_i values (right) obtained by the network. (b) Imaginary part of the characteristic function for the data in Figure 1 (left) and the mesh of θ_i values (right) obtained by the network.

$$\chi_{(ij)}(\vec{\alpha}) = \frac{e^{-|\vec{\alpha} - \vec{\alpha}_{(ij)}|^2 / 2\sigma_{ij}^2}}{\sum_{(kl)} e^{-|\vec{\alpha} - \vec{\alpha}_{(kl)}|^2 / 2\sigma_{kl}^2}} \quad (10)$$

Two networks are used, one for the real part of the characteristic function, another for the imaginary part with, in eqn (9) $\cos(\vec{\alpha}_{(ij)} \cdot \vec{x}(t))$ replaced by $\sin(\vec{\alpha}_{(ij)} \cdot \vec{x}(t))$.

Figure 5(a,b) shows the values computed by our algorithm for the real and imaginary parts of the characteristic function corresponding to the two-dimensional signal in Figure 1. On the left is a plot of the exact characteristic function and on the right the values learned by the network. In this case we show only the mesh corresponding to the θ_i values. One obtains a 2.0% accuracy for the real part and 4.5% accuracy for the imaginary part.

The convergence of the learning process is fast and the approximation is reasonably good. Notice in particular the slope discontinuity at the origin which reveals the non-existence of a second moment. The parameters used for the learning laws in this example were $\gamma = 0.00036$, $\eta = 1.8$, $\sigma = 0.25$. The number of points in the training set is 25,000.

For a second example the data was generated by a Weierstrass random walk with probability distribution

$$p(x) = \frac{1}{6} \sum_{j=0}^{\infty} \left(\frac{2}{3}\right)^j (\delta_{x, b^j} + \delta_{x, -b^j}) \quad (11)$$

and $b = 1.31$, which is a process of the Lévy flight type. The characteristic function, obtained by the network, is shown in Figure 6. Taking the $\log(-\log)$ of the network output one obtains the scaling exponent 1.49 near $\alpha = 0$, close to the expected fractal dimension of the random walk path (1.5). The parameters used for the learning laws in this example were $\gamma = 0.0005$, $\eta = 1.75$, $\sigma = 0.1732$. The number of points in the training set is 80,000.

These examples test the algorithm as a process identifier, in the sense that, after the learning process, the network is a dynamical representation of the characteristic function and may be used to perform all kinds of analysis of the statistics of the data.

There are other ways to obtain the characteristic function of a probability distribution, which may be found in the statistical inference literature (Prakasa Rao, 1987). Our purpose in developing neural-like algorithms for this purpose was both to have a device that, after learning, is quick to evaluate and, at the same time, adjusts itself easily, through continuous learning, to changing statistics. As the PCA algorithms that extract the full correlation matrix, our neural algorithm laws are also non-local. As a computation algorithm this is not a serious issue, but for hardware implementations it might raise some problems.




FIGURE 6. (a) Characteristic function for the Weierstrass random walk ($b = 1.31$). $\text{Log}(-\log)$ of the characteristic function $f(a)$ for the Weierstrass random walk ($b = 1.31$).

3. IDENTIFICATION OF STOCHASTIC PROCESSES

As we have stated before the full characterisation of the time structure of a stochastic process requires the knowledge of its characteristic functional (eqn (3)) for a dense set of functions $\xi(t)$.

To construct an approximation to the characteristic functional we might discretise the time $\{tm = n\Delta t\}$ and the inner product in the exponential becomes a sum over the process sampled at a sequence of times.

$$F(\xi) = \left\langle e^{i \sum_k x(t_k)\xi(t_k)} \right\rangle \quad (12)$$

The problem would then be reduced to the construction of a multidimensional characteristic function as in Section 2. In practice we would have to limit the time-depth of the functional to a maximum of T time steps, $T\Delta t$ being the maximum time delay over which time correlations may be obtained. If N is the number of different ξ values for each k , the algorithm developed in Section 2 requires N^T nodes in the intermediate layer

and, for any reasonably large T , this method becomes computationally explosive.

An alternative and computationally simpler method consist in, restricting ourselves to Markov processes, to characterise the process by the construction of networks to represent the transition function for fixed time intervals. From these networks the differential Chapman-Kolmogorov equation may then be reconstructed.

Let $x(t)$ be a one-dimension Markov process and $p(x_2, t + \Delta t | x_1, t)$ its transition function that is, the conditional probability of finding the value x_2 at time $t + \Delta t$ given x_1 at time t . Assume further that the process is stationary

$$p(x_2, t + \Delta t | x_1, t) = p(x_2, \Delta t | x_1). \tag{13}$$

The network that configures itself to represent this function is similar to the one we used for the two-dimensional characteristic function. It is sketched in Figure 7(a,b). It has a set of N^2 intermediate layer nodes with node parameters, the node with coordinates $\vec{x}_{(ij)}$ corresponding to the arguments $(x_2(ij) = x_0 + i\Delta x, x_1(ij) = x_0 + j\Delta x)$ in the transition function. The domain of both arguments is $(x_0, x_0 + N\Delta x)$. For each pair $\{x_2 = x(t + \Delta t), x_1 = x(t)\}$ in the data set, the node parameters that are updated are those in the four columns containing the nearest neighbours of the point $\vec{x} = \{x_2, x_1\}$ (Figure 7b).

The learning rule is

$$\theta_{(ij)}(t+1) = \frac{S_j(t)\theta_{(ij)}(t) + \frac{\eta_{(ij)}(\vec{x})e^{-|\vec{x}-\vec{x}_{(ij)}|^2/\alpha}}{\sum_{kl} \eta_{(kl)}(\vec{x})e^{-|\vec{x}-\vec{x}_{(kl)}|^2/\alpha}}}{S_j(t+1)} \tag{14}$$

$$S_j(t+1) = S_j(t) + \sum_i \frac{\eta_{(ij)}(\vec{x})e^{-|\vec{x}-\vec{x}_{(ij)}|^2/\alpha}}{\sum_{kl} \eta_{(kl)}(\vec{x})e^{-|\vec{x}-\vec{x}_{(kl)}|^2/\alpha}} \tag{15}$$

where $\eta_{ij}(\vec{x}) = 1$ if (ij) is one of the nearest neighbours of the data point and zero otherwise. α is a neighbourhood smoothing factor. $S_j(t)$ is a column normalisation factor.




FIGURE 7. (a) Network that learns the transition function $p(x_2, \Delta t | x_1)$ of a stationary Markov process; (b) Nearest neighbours of the data point $\vec{x} = (x(t + \Delta t), x(t))$.




FIGURE 8. Typical sample path of a stationary Markov process.

In the limit of large learning times the node parameters approach the transition function

$$\theta_{(ij)} \rightarrow p(x_0 + i\Delta x, \Delta t | x_0 = j\Delta x). \tag{16}$$

As for the networks in Section 2, the intermediate layer nodes are equipped with a radial basis function [eqn (10)] and the connection strengths in the output additive node have a learning law identical to the second equation in eqn (9). The role of this part of the network is, as before, to obtain an interpolating effect.

What the algorithm of eqns (14) and (15) does is to compute recursively the average number of transitions between points in the configuration space of the process. The spatial smoothing effect of the algorithm automatically insures a good representation of a continuous function from a finite data set. Furthermore its recursive nature would be appropriate for the case of drifting statistics.

For a stationary process, once the learning process has been achieved and if Δt is chosen to be sufficiently small, the network itself may be used to simulate the stationary Markov process. A complete characterisation of the process may also be obtained by training a few similar networks for different (small) Δt values and computing the coefficient functions in the differential Chapman-Kolmogorov equation (Gardiner, 1983).




FIGURE 9. Transition functions obtained for (a) $t = 3\Delta T$ and (b) $t = \Delta T$.




FIGURE 10. Transition function $p(x, \Delta T|0, 0)$ (a) and $p(x, 3\Delta T|0, 0)$ (b).


$$\begin{aligned} \partial_t p(\vec{z}, t'|\vec{y}, t) = & - \sum_i \frac{\partial}{\partial z_i} \{A_i(\vec{z}, t') p(\vec{z}, t'|\vec{y}, t)\} \\ & + \sum_{ij} \frac{1}{2} \frac{\partial^2}{\partial z_i \partial z_j} \{B_{ij}(\vec{z}, t') p(\vec{z}, t'|\vec{y}, t)\} \\ & + \int d\vec{x} \{W(\vec{z}|\vec{x}, t') p(\vec{x}, t'|\vec{y}, t) \\ & - W(\vec{x}|\vec{z}, t') p(\vec{z}, t'|\vec{y}, t)\}. \end{aligned} \quad (17)$$

The coefficients are obtained from the transition probabilities, noticing that for all $\epsilon > 0$.

$$\lim_{\Delta t \rightarrow 0} p(\vec{x}, t + \Delta t|\vec{z}, t) / \Delta t = W(\vec{x}|\vec{z}, t) \text{ for } |\vec{x} - \vec{z}| \geq \epsilon \quad (18)$$

$$\lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int_{|\vec{x} - \vec{z}| < \epsilon} dx (x_i - z_i) p(\vec{x}, t + \Delta t|\vec{z}, t) = A_i(\vec{z}, t) + O(\epsilon) \quad (19)$$

$$\begin{aligned} \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int_{|\vec{x} - \vec{z}| < \epsilon} dx (x_i - z_i)(x_j - z_j) p(\vec{x}, t + \Delta t|\vec{z}, t) \\ = B_{ij}(\vec{z}, t) + O(\epsilon). \end{aligned} \quad (20)$$



(a)





FIGURE 12. Diffusion function for ΔT (o), $2\Delta T$ (•) and $3\Delta T$ (▪).

$W(\vec{x}|\vec{z}, t)$ is the jumping kernel, $A_i(\vec{z}, t)$ the drift and $B_{ij}(\vec{z}, t)$ the diffusion coefficient.

As an example we have considered a Markov process with jumping, drift and diffusion. A typical sample path is shown in Figure 8. Three networks were trained on this process, to learn the transition function for $t = \Delta T, 2\Delta T$ and $3\Delta T$ ($\Delta T = 0.374$ ms). Figure 9 shows the transition function for $t = \Delta T$ and $3\Delta T$. Figure 10 shows two sections of the transition function for $x_2 = 0$, that is $p(x, \Delta T|0, 0)$ and $p(x, 3\Delta T|0, 0)$.

The networks were then used to extract the coefficient functions $A(x, t)$, $B(x, t)$ and $W(x|z, t)$. To find the drift $A(x, t)$ we use eqn (19). Figure 11 shows the computed drift function and a least square linear fit. Also shown is a comparison with the exact drift function of the process.

To obtain the diffusion coefficient $B(x, t)$ we use eqn (20). Figure 12 shows the diffusion coefficient for different ΔT values. ΔT is the smallest time step used in the process simulation. Therefore $B(x, t) = 2.6$ is our estimate for the diffusion coefficient. In this case, because the diffusion coefficient is found to be a constant, the value of the jumping kernel $W(x|z, t)$ is easily obtained by integration around the local maxima x_m of



(b)

FIGURE 11. (a) Computed drift function and linear least square fit; (b) exact drift function of the process (1), least square linear fits for (2) $t = \Delta T$ and (3) $t = 2\Delta T$.

$p(x, \Delta T|z)$ with $|x - z| > 0.2$.

$$W = \frac{1}{\Delta T} \int_{x_m - \delta}^{x_m + \delta} p(x, \Delta T|z) dx \quad \delta = 0.2. \quad (21)$$

We conclude $W(x, z) \approx 300\delta(|x - z| - 0.5)$. The parameters used for the learning laws in this example were $\eta = 0.48$, $\alpha = 0.00021$. The number of points in the training set is 1,000,000.

REFERENCES

- Dente, J. A., & Vilela Mendes, R. (1996). Unsupervised learning in general connectionist systems. *Network: Computation in Neural Systems*, 7, 123.
- Doyle Farmer, J. (1990). A Rosetta stone for connectionism. *Physica*, D42, 153.
- Gardiner, C. W. (1983). *Handbook of stochastic methods*. Berlin: Springer.
- Hebb, D. O. (1949). *The organisation of behaviour*. New York: Wiley.
- Hida, T. (1980). *Brownian motion*. Berlin: Springer.
- Lukacs, E. (1970). *Characteristic functions*. London: Griffin.
- Oja, E. (1989). Neural networks, principal components and subspaces. *International Journal of Neural Systems*, 1, 61.
- Oja, E. (1992). Principal components, minor components and linear neural networks. *Neural Networks*, 5, 927.
- Prakasa Rao, B. L. S. (1987). *Asymptotic theory of statistical inference*. New York: Wiley.
- Sanger, T. D. (1989). Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks*, 2, 459.
- Shlesinger, M. F., Zaslavski, G. M., & Klafter, J. (1993). Strange kinetics. *Nature*, 363, 31.
- Softky, W. R., & Kammen, D. M. (1991). Correlations in high dimensional or asymmetric data sets: Hebbian neuronal processing. *Neural Networks*, 4, 337.
- Taylor, J. G., & Coombes, S. (1993). Learning higher order correlations. *Neural Networks*, 6, 423.
- Zumofen, G., & Klafter, J. (1993). Scale-invariant motion in intermittent chaotic systems. *Physics Review*, E47, 851.
- Zumofen, G., & Klafter, J. (1994). Random walks in the standard map. *Europhysics Letters*, 25, 565