Estimating probability densities from short samples: a parametric maximum likelihood approach

Thierry Dudok de Wit, E. Floriani

► To cite this version:

Thierry Dudok de Wit, E. Floriani. Estimating probability densities from short samples: a parametric maximum likelihood approach. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, 1998, 58 (4), pp.5115. 10.1103/PhysRevE.58.5115 . hal-00823527

HAL Id: hal-00823527

<https://hal.science/hal-00823527>

Submitted on 12 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Estimating probability densities from short samples: A parametric maximum likelihood approach

T. Dudok de Wit* and E. Floriani†

Centre de Physique Théorique, Luminy Case 907, 13288 Marseille Cedex 9, France

(Received 16 January 1998)

A parametric method similar to autoregressive spectral estimators is proposed to determine the probability density function (PDF) of a random set. The method proceeds by maximizing the likelihood of the PDF, yielding estimates that perform equally well in the tails as in the bulk of the distribution. It is therefore well suited for the analysis of short sets drawn from smooth PDF's and stands out by the simplicity of its computational scheme. Its advantages and limitations are discussed. [S1063-651X(98)06410-1]

PACS number(s): 02.70.Hm, 02.50.Ng

I. INTRODUCTION

There are many applications in which it is necessary to estimate the probability density function PDF from a finite sample of n observations $\{x_1, x_2, \dots, x_n\}$ whose true PDF is $f(x)$. Here we consider the generic case in which the identically distributed (but not necessarily independent) random variables have a compact support $x_k \in [a, b]$.

The usual starting point for a PDF estimation is the naive estimate

$$\hat{f}_\delta(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - x_i), \quad (1)$$

where $\delta(\cdot)$ stands for the Dirac delta function. Although this definition has a number of advantages, it is useless for practical purposes since a smooth functional is needed.

Our problem consists in finding an estimate $\hat{f}(x)$ whose integral over an interval of given length converges toward that of the true PDF as $n \rightarrow \infty$. Many solutions have been developed for that purpose: foremost among these are kernel techniques in which the estimate $\hat{f}_\delta(x)$ is smoothed locally using a kernel function $K(x)$ [1–3],

$$\hat{f}(x) = \int_a^b \frac{1}{w} K\left(\frac{x-y}{w}\right) \hat{f}_\delta(y) dy, \quad (2)$$

whose width is controlled by the parameter w . The well-known histogram is a variant of this technique. Although kernel approaches are by far the most popular ones, the choice of a suitable width remains a basic problem for which visual guidance is often needed. More generally, one faces the problem of choosing a good partition. Some solutions include Bayesian approaches [4], polynomial fits [5], and methods based on wavelet filtering [6].

An alternative approach, considered by many authors [7–11], is a projection of the PDF on orthogonal functions

$$\hat{f}(x) = \sum_k \alpha_k g_k(x), \quad (3)$$

where the partition problem is now treated in dual space. This parametric approach has a number of interesting properties: a finite expansion often suffices to obtain a good approximation of the PDF and the convergence of the series versus the sample size n is generally faster than for kernel estimates. A strong point is its global character, since the PDF is fitted globally, yielding estimates that are better behaved in regions where the lack of statistics causes kernel estimates to perform poorly. Such a property is particularly relevant for the analysis of turbulent wave fields, in which the tails of the distribution are of great interest (e.g., [12]).

These advantages, however, should be weighed against a number of downsides. Orthogonal series do not provide consistent estimates of the PDF since for increasing number of terms they converge toward $\hat{f}_\delta(x)$ instead of the true density $f(x)$ [13]. Furthermore, most series can only handle continuous or piecewise continuous densities. Finally, the PDF estimates obtained that way are not guaranteed to be nonnegative (see, for example, the problems encountered in [14]).

The first problem is not a major obstacle, since most experimental distributions are smooth anyway. The second one is more problematic. In this paper we show how it can be partly overcome by using a Fourier series expansion of the PDF and seeking a maximization of the likelihood

$$\hat{L} = \int_a^b \ln \hat{f}(x) dx. \quad (4)$$

The problem of choosing an appropriate partition then reduces to that of fitting the PDF with a positive definite Padé approximant [15].

Our motivation for presenting this particular parametric approach stems from its robustness, its simplicity, and the originality of the computational scheme it leads to. The latter, as will be shown later, is closely related to the problem of estimating power spectral densities with autoregressive (AR) or maximum entropy methods [16–18]. To the best of our knowledge, the only earlier reference to similar work is that by Carmichael [19]; here we emphasize the relevance of

*Affiliated with: Université de Provence and CNRS. Electronic address: ddwit@cpt.univ.mrs.fr

†Electronic address: floriani@cpt.univ.mrs.fr

the method for estimating PDF's and propose a criterion for choosing the optimum number of basis functions.

II. THE MAXIMUM LIKELIHOOD APPROACH

The method we now describe basically involves a projection of the PDF on a Fourier series. The correspondence between the continuous PDF $f(x)$ and its discrete characteristic function ϕ_k is established by [20]

$$\phi_k = \int_{-\pi}^{+\pi} f(x) e^{jkx} dx, \quad (5)$$

$$f(x) = 2\pi \sum_{k=-\infty}^{+\infty} \phi_k e^{-jkx}, \quad (6)$$

where $\phi_k = \phi_{-k}^* \in \mathbb{C}$ is Hermitian [21]. Note that we have applied a linear transformation to convert the support from $[a, b]$ to $[-\pi, \pi]$.

For a finite sample, an unbiased estimate of the characteristic function is obtained by inserting Eq. (1) into Eq. (5), giving

$$\hat{\phi}_k = \frac{1}{n} \sum_{i=1}^n e^{jkx_i}. \quad (7)$$

The main problem now consists in recovering the PDF from Eq. (6) while avoiding the infinite summation. By working in dual space we have substituted the partition choice problem by that of selecting the number of relevant terms in the Fourier series expansion.

The simplest choice would be to truncate the series at a given "wave number" p and discard the other ones

$$\hat{f}(x) = 2\pi \sum_{k=-p}^{+p} \hat{\phi}_k e^{-jkx}. \quad (8)$$

Such a truncation is equivalent to keeping the lowest wave numbers and thus filtering out small details of the PDF. Incidentally, this solution is equivalent to a kernel filtering with $K(x) = \sin(\pi x)/\pi x$ as kernel. This kernel is usually avoided because it suffers from many drawbacks such as the generation of spurious oscillations.

An interesting improvement was suggested by Burg in the context of spectral density estimation (see, for example, [16,17]). The heuristic idea is to keep some of the low wave number terms while the remaining ones, instead of being set to zero, are left as free parameters:

$$\hat{f}(x) = 2\pi \sum_{k=-\infty}^{+\infty} \hat{\alpha}_k e^{-jkx} \quad \text{with} \quad \hat{\alpha}_k = \hat{\phi}_k, \quad |k| \leq p. \quad (9)$$

The parameters $\hat{\alpha}_k$, for $|k| > p$, are then fixed self-consistently according to some criterion.

We make use of this freedom to constrain the solution to a particular class of estimates. Without any prior information at hand, a reasonable choice is to select the estimate that contains the least possible information or is the most likely.

It is therefore natural to seek a maximization of an entropic quantity such as the sample entropy

$$\hat{H} = - \int_{-\pi}^{+\pi} \hat{f}(x) \ln \hat{f}(x) dx, \quad (10)$$

or the sample likelihood

$$\hat{L} = \int_{-\pi}^{+\pi} \ln \hat{f}(x) dx. \quad (11)$$

We are *a priori* inclined to choose the entropy because our objective is the estimation of the PDF and not that of the characteristic function. However, numerical investigations done in the context of spectral density estimation rather lend support to the likelihood criterion [22]. A different and stronger motivation for preferring a maximization of the likelihood comes from the simplicity of the computational scheme it gives rise to.

This maximization means that the tail of the characteristic function is chosen subject to the constraint

$$\frac{\partial \hat{L}}{\partial \hat{\alpha}_k} = 0, \quad |k| > p. \quad (12)$$

From Eqs. (9) and (11) the likelihood can be rewritten as

$$\hat{L} = \int_{-\pi}^{+\pi} \ln \left(2\pi \sum_{k=-\infty}^{+\infty} \hat{\alpha}_k e^{-jkx} \right) dx. \quad (13)$$

As shown in the Appendix, the likelihood is maximized when the PDF can be expressed by the functional

$$\hat{f}_p(x) = \frac{1}{\sum_{k=-p}^p c_k e^{-jkx}}, \quad (14)$$

which is a particular case of a Padé approximant with poles only and no zeros [15]. Requiring that $\hat{f}_p(x)$ is real and bounded, it can be rewritten as

$$\hat{f}_p(x) = \frac{\varepsilon_0}{2\pi} \frac{1}{|1 + a_1 e^{-jx} + \dots + a_p e^{-jpx}|^2}. \quad (15)$$

The values of the coefficients $\{a_1, \dots, a_p\}$ and of the normalization constant ε_0 are set by the condition that the Fourier transform of $\hat{f}_p(x)$ must match the sample characteristic function $\hat{\phi}_k$ for $|k| \leq p$.

This solution has a number of remarkable properties, some of which are deferred to the Appendix. Foremost among these are its positive definite character and the simple relationship which links the polynomial coefficients $\{a_1, \dots, a_p\}$ to the characteristic function on which they perform a regression. Indeed, we have

$$\hat{\phi}_k + a_1 \hat{\phi}_{k-1} + a_2 \hat{\phi}_{k-2} + \dots + a_p \hat{\phi}_{k-p} = 0, \quad 1 \leq k \leq p. \quad (16)$$

This can be cast in a set of Yule-Walker equations whose unique solution contains the polynomial coefficients

$$\begin{bmatrix} \phi_0 & \phi_{-1} & \cdots & \phi_{-p+1} \\ \phi_1 & \phi_0 & \cdots & \phi_{-p+2} \\ \vdots & \vdots & & \vdots \\ \phi_{p-1} & \phi_{p-2} & \cdots & \phi_0 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_p \end{bmatrix} = - \begin{bmatrix} \phi_1 \\ \phi_2 \\ \vdots \\ \phi_p \end{bmatrix}. \quad (17)$$

Advantage can be taken here of the Toeplitz structure of the matrix. The proper normalization $[\int_{-\pi}^{+\pi} \hat{f}(x) dx = 1]$ of the PDF is ensured by the value of ε_0 , which is given by a variant of Eq. (16),

$$\phi_0 + a_1 \phi_{-1} + a_2 \phi_{-2} + \cdots + a_p \phi_{-p} = \varepsilon_0. \quad (18)$$

Equations (15) and (17) illustrate the simplicity of the method.

III. SOME PROPERTIES

A clear advantage of the method over conventional series expansions is the automatic positive definite character of the PDF. Another asset is the close resemblance with autoregressive or maximum entropy methods that are nowadays widely used in the estimation of spectral densities. Both methods have in common the estimation of a positive function by means of a Padé approximant whose coefficients directly issue from a regression [Eq. (16)]. This analogy allows us to exploit here some results previously obtained in the framework of spectral analysis.

One of these concerns the statistical properties of the maximum likelihood estimate. These properties are badly known because the nonlinearity of the problem impedes any analytical treatment. The analogy with spectral densities, however, reveals that the estimates are asymptotically normally distributed with a standard deviation [23,24]

$$\sigma_{\hat{f}} \propto \sqrt{\hat{f}}. \quad (19)$$

This scaling should be compared against that of conventional kernel estimates, for which

$$\sigma_{\hat{f}} \propto \sqrt{\hat{f}}. \quad (20)$$

The key point is that kernel estimates are relatively less reliable in low density regions than in the bulk of the distribution, whereas the relative uncertainty of maximum likelihood estimates is essentially constant. The latter property is obviously preferable when the tails of the distribution must be investigated, e.g., in the study of rare events.

Some comments are now in order. By choosing a Fourier series expansion, we have implicitly assumed that the PDF was 2π periodic, which is not necessarily the case. Thus special care is needed to enforce periodicity, since otherwise wraparound may result [25]. The solution to this problem depends on how easily the PDF can be extended periodically. In most applications, the tails of the distribution progressively decrease to zero, so periodicity may be enforced simply by artificially padding the tails with a small interval in which the density vanishes. We do this by rescaling the support from $[a, b]$ to an interval which is slightly smaller

than 2π , say $[-3, 3]$ [26]. Once the Padé approximant is known, the $[-3, 3]$ interval is scaled back to $[a, b]$.

If there is no natural periodic extension to the PDF [for example, if $f(a)$ strongly differs from $f(b)$] then the choice of Fourier basis functions in Eq. (3) becomes questionable and, not surprisingly, the quality of the fit degrades. Even in this case, however, the results can still be improved by using *ad hoc* solutions [27].

We mentioned before that the maximum likelihood method stands out by its computational simplicity. Indeed, a minimization of the entropy would lead to the solution

$$\ln \hat{f}_p(x) \propto \sum_{k=-p}^p c_k e^{-jkx}, \quad (21)$$

whose numerical implementation requires an iterative minimization and is therefore considerably more demanding.

Finally, the computational cost is found to be comparable to or even better (for large sets) than for conventional histogram estimates. Most of the computation time goes into the calculation of the characteristic function, for which the number of operations scales as the sample size n .

IV. CHOOSING THE ORDER OF THE MODEL

The larger the order p of the model is, the finer the details in the PDF estimate are. Finite sample effects, however, also increase with p . It is therefore of prime importance to find a compromise. Conventional criteria for selecting the best compromise between model complexity and quality of the fit, such as the final prediction error and the minimum description length [16–18], are not applicable here because they require the series of characteristic functions $\{\phi_k\}$ to be normally distributed, which they are not.

Guided by the way these empirical criteria have been chosen, we have defined a new one, which is based on the following observation: as p increases starting from 0, the PDF's $\hat{f}_p(x)$ progressively converge toward a stationary shape; after some optimal order, however, ripples appear and the shapes start diverging again. It is therefore reasonable to compare the PDF's pairwise and determine how close they are. A natural measure of closeness between two positive distributions $\hat{f}_p(x)$ and $\hat{f}_{p+1}(x)$ is the Kullback-Leibler entropy or information gain [28,29]

$$\hat{I}(\hat{f}_{p+1}, \hat{f}_p) = \int_{-\pi}^{+\pi} \hat{f}_{p+1}(x) \ln \frac{\hat{f}_{p+1}(x)}{\hat{f}_p(x)} dx, \quad (22)$$

which quantifies the amount of information gained by changing the probability density describing our sample from $\hat{f}_p$ to $\hat{f}_{p+1}$. In other words, if H_p (or H_{p+1}) is the hypothesis that x was selected from the population whose probability density is $\hat{f}_p$ ($\hat{f}_{p+1}$), then $\hat{I}(\hat{f}_{p+1}, \hat{f}_p)$ is given as the mean information for discriminating between H_{p+1} and H_p per observation from $\hat{f}_{p+1}$ [28].

Notice that the information gain is not a distance between distributions; it nevertheless has the property of being non-negative and to vanish if and only if $\hat{f}_p \equiv \hat{f}_{p+1}$. We now proceed as follows: starting from $p=0$ the order is incre-

mented until the information gain reaches a clear minimum; this corresponds, as it has been checked numerically, to the convergence toward a stationary shape; the corresponding order is then taken as the requested compromise. Clearly, there is some arbitrariness in the definition of such a minimum since visual inspection and common sense are needed. In most cases, however, the solution is evident and the search can be automated. Optimal orders usually range between 2 and 10; larger values may be needed to model discontinuous or complex shaped densities.

V. SOME EXAMPLES

Three examples are now given in order to illustrate the limits and the advantages of the method.

A. General properties

First, we consider a normal distribution with exponential tails as often encountered in turbulent wave fields. We simulated a random sample with $n = 2000$ elements and the main results appear in Fig. 1.

The information gain [Fig. 1(b)] decreases as expected until it reaches a well-defined minimum at $p = 7$, which therefore sets the optimal order of our model. Since the true PDF is known, we can test this result against a common measure of the quality of the fit, which is the mean integrated squared error (MISE), d ,

$$d(p) = \int_a^b [f(x) - \hat{f}_p(x)]^2 dx. \quad (23)$$

The MISE, which is displayed in Fig. 1(b), also reaches a minimum at $p = 7$ and thus supports the choice of the information gain as a reliable indicator for the best model. Tests carried out on other types of distributions confirm this good agreement.

Now that the optimum PDF has been found, its characteristic function can be computed and compared with the measured one, see Fig. 1(a). As expected, the two characteristic functions coincide for the p lowest wave numbers [Eq. (16)]; they diverge at higher wave numbers, for which the model tries to extrapolate the characteristic function self-consistently. The fast falloff of the maximum likelihood estimate explains the relatively smooth shape of the resulting PDF.

Finally, the quality of the PDF can be visualized in Fig. 1(d), which compares the measured PDF with the true one, and an estimate based on a histogram with 101 bins. An excellent agreement is obtained, both in the bulk of the distribution and in the tails, where the exponential falloff is correctly reproduced. This example illustrates the ability of the method to get reliable estimates in regions where standard histogram approaches have a lower performance.

B. Interpreting the characteristic function

The shape of the characteristic function in Fig. 1(a) is reminiscent of spectral densities consisting of a low wave number (band-limited) component embedded in broadband

FIG. 1. Example of a normal distribution with exponential tails. The sample size is $n = 2000$. From top to bottom are shown: (a) the magnitude $|\phi_k|$ of the characteristic function (thick line) and the fit resulting from a seventh order model; (b) the information gain (thick line) and the MISE, both showing a minimum around $p = 7$ which is marked by a circle; (c) the likelihood $\hat{L}$ associated with the different PDF's estimated for $p = 1 - 20$; and finally (d) the maximum likelihood estimate of the PDF (thick line), an estimate based on a histogram with 101 equispaced bins (dots) and the true PDF (thin line).

noise. A straightforward calculation of the expectation of $|\phi_k|$ indeed reveals the presence of a bias which is due to the finite sample size

$$E[|\phi_k|] = |\phi_k| + \frac{\gamma}{\sqrt{n}}, \quad (24)$$

where γ depends on the degree of independence between the samples in $\{x\}$. This bias is illustrated in Fig. 2 for independent variables drawn from a normal distribution, showing how the wave number resolution gradually degrades as the sample size decreases. Incidentally, a knowledge of the bias level could be used to obtain confidence intervals for the PDF estimate. This would be interesting insofar as no assumptions have to be made on possible correlations in the data set. We found this approach, however, to be too inaccurate on average to be useful.

FIG. 2. The expectation $E[|\hat{\phi}_k|]$ computed for sets of various sizes taken from the same normal distribution. The noise-induced bias level goes down as the size increases, progressively revealing finer details of the PDF.

The presence of a bias also gives an indication of the smallest scales (in terms of amplitude of x) one can reliably distinguish in the PDF. For a set of 2000 samples drawn from a normal distribution, for example, components with wave numbers in excess of $k=3$ are hidden by noise and hence the smallest meaningful scales in the PDF are of the order of $\delta x=0.33$. These results could possibly be further improved by Wiener filtering.

C. Influence of the sample size

To investigate the effect of the sample length n , we now consider a bimodal distribution consisting of two normal distributions with different means and standard deviations. Such distributions are known to be difficult to handle with kernel estimators.

Samples with, respectively, $n=200$, $n=2000$, and $n=20\,000$ elements were generated; their characteristic functions and the resulting PDF's are displayed in Fig. 3. Clearly, finite sample effects cannot be avoided for small samples but the method nevertheless succeeds relatively well in capturing the true PDF and in particular the small peak associated with the narrow distribution. An analysis of the MISE shows that it is systematically lower for maximum likelihood estimates than for standard histogram estimates, supporting the former.

D. A counterexample

The previous examples gave relatively good results because the true distributions were rather smooth. Although such smooth distributions are generic in most applications it may be instructive to look at a counterexample, in which the method fails.

Consider the distribution which corresponds to a cut through an annulus

$$f(x) = \begin{cases} \frac{1}{2}, & 1 \leq |x| \leq 2 \\ 0 & \text{elsewhere.} \end{cases} \quad (25)$$

A sample was generated with $n=2000$ elements and the resulting information gains are shown in Fig. 4. There is an ambiguity in the choice of the model order and indeed the convergence of the PDF estimates toward the true PDF is neither uniform nor in the mean. Increasing the order improves the fit of the discontinuity a little but also increases the oscillatory behavior known as the Gibbs phenomenon.

FIG. 3. The PDF's as calculated for sets of various sizes taken from the same binormal distribution. The thick line designates the maximum likelihood estimate, the thin line the true PDF, and the dots a histogram estimate obtained from 61 equispaced bins. The optimum orders are, respectively, from top to bottom $p=5$, $p=6$, and $p=11$.

This problem is related to the fact that the PDF is discontinuous and hence the characteristic function is not absolutely summable.

Similar problems are routinely encountered in the design of digital filters, where steep responses cannot be approximated with infinite impulse response filters that have a limited number of poles [20]. The bad performance of the maximum likelihood approach in this case also comes from its inability to handle densities that vanish over finite intervals. A minimization of the entropy would be more appropriate here.

VI. CONCLUSION

We have presented a parametric procedure for estimating univariate densities using a positive definite functional. The method proceeds by maximizing the likelihood of the PDF subject to the constraint that the characteristic functions of the sample and estimated PDF's coincide for a given number of terms. Such a global approach to the estimation of PDF's is in contrast to the better known local methods (such as nonparametric kernel methods) whose performance is poorer in regions where there is a lack of statistics, such as the tails of the distribution. This difference makes the maximum likelihood method relevant for the analysis of short records (with typically hundreds or thousands of samples). Other advantages include a simple computational procedure that can be tuned with a single parameter. An entropy-based criterion

FIG. 4. Results obtained for an annular distribution; the sample size is $n=2000$. In (a) the information gain has no clear minimum and hence there is no well-defined order for the model. In (b) the estimated PDF's for $p=1$ and $p=2$ fail to fit the true PDF (thick line). Increasing the order (c) improves the fit but also enhances the Gibbs phenomenon. Dots correspond to a histogram estimate with equispaced bins.

has been developed for selecting the latter.

The method works best with densities that are at least once continuously differentiable and that can be extended periodically. Indeed, the shortcomings of the method are essentially the same as for autoregressive spectral estimates, which give rise to the Gibbs phenomenon if the density is discontinuous.

The method can be extended to multivariate densities, but the computational procedures are not yet within the realm of practical usage. Its numerous analogies with the design of digital filters suggest that it is still open to improvements.

ACKNOWLEDGMENTS

We gratefully acknowledge the dynamical systems team at the Center de Physique Théorique for many discussions as well as D. Lagoutte and B. Torrèsani for making comments on the manuscript. E. Floriani acknowledges support by the EC under Grant No. ERBFMBICT960891.

APPENDIX

We detail here the main stages that lead to the PDF estimate described in Sec. II because extensive proofs are rather

difficult to find in the literature.

The maximum likelihood condition [Eq. (12)] can be expressed as

$$\int_{-\pi}^{+\pi} \frac{e^{-jkx}}{\sum_{l=-\infty}^{\infty} \hat{\alpha}_l e^{-jl x}} dx = \int_{-\pi}^{+\pi} \frac{e^{-jkx}}{\hat{f}(x)} dx = 0, \quad (\text{A1})$$

for $|k| > p$ [30]. This simply means that the Fourier expansion of $[\hat{f}(x)]^{-1}$ should not contain terms of order $|k| > p$ and hence the solution must be

$$\hat{f}_p(x) = \frac{1}{\sum_{k=-p}^p c_k e^{-jkx}}. \quad (\text{A2})$$

The PDF we are looking for must of course be real, and so the coefficients should be Hermitian $c_k = c_{-k}^*$. We also want the PDF to be bounded, which implies that

$$\sum_{k=-p}^p c_k e^{-jkx} \neq 0, \quad \forall x \in [-\pi, \pi]. \quad (\text{A3})$$

Let us now define, for z complex

$$C(z) = \sum_{k=-p}^p c_{-k} z^k, \quad (\text{A4})$$

and

$$P(z) = z^p C(z). \quad (\text{A5})$$

$P(z)$ is a polynomial of degree $2p$. It can be easily verified that [31]

$$P(z) = z^{2p} \left[P\left(\frac{1}{z^*}\right) \right]^* \quad (\text{A6})$$

as a consequence of the Hermiticity of the coefficients c_k . In particular, this tells us that if z_1 is a root of $P(z)$, then $1/z_1^*$ (the complex conjugate of its mirror image with respect to the unit circle) is also a root of $P(z)$. From Eq. (A3) we know that none of these roots are located on the unit circle.

Let us now rearrange the roots of $P(z)$, denoting by $\{z_1, \dots, z_p\}$ the p roots lying outside the unit disk and by $\{1/z_1^*, \dots, 1/z_p^*\}$ the p other ones that are located inside the unit circle. We can then write

$$P(z) = c_{-p}(z - z_1) \cdots (z - z_p) \left(z - \frac{1}{z_1^*} \right) \cdots \left(z - \frac{1}{z_p^*} \right), \quad (\text{A7})$$

with

$$c_{-p} z_1 \cdots z_p = c_p z_1^* \cdots z_p^*. \quad (\text{A8})$$

From this $C(z)$ can be written as

$$C(z) = \pm B(z) \left[B\left(\frac{1}{z^*}\right) \right]^*, \quad (\text{A9})$$

where

$$B(z) = \left| \frac{c_p}{z_1 \cdots z_p} \right|^{1/2} (z - z_1) \cdots (z - z_p). \quad (\text{A10})$$

By construction, all the roots of $B(z)$ are located outside the unit disk.

Finally, we get for $\hat{f}_p(x)$

$$\hat{f}_p(x) = \frac{1}{C(z=e^{jx})} = \pm \frac{1}{|B(e^{jx})|^2}. \quad (\text{A11})$$

All the solutions of the maximum likelihood principle, if real and bounded, are thus of constant sign and have the structure given by Eq. (A11). Excluding negative definite solutions we obtain

$$\hat{f}_p(x) = \frac{\varepsilon_0}{2\pi} \frac{1}{|1 + a_1 e^{-jx} + \cdots + a_p e^{-jpx}|^2}, \quad (\text{A12})$$

where

$$\varepsilon_0 = \frac{2\pi}{|b_0|^2}, \quad a_i = \frac{b_i^*}{b_0^*}, \quad i = 1, \dots, p, \quad (\text{A13})$$

where $\{b_1, \dots, b_p\}$ are the coefficients of the polynomial $B(z)$ and $A(z) = 1 + a_1 z + \cdots + a_p z^p$ has all its roots outside the unit disk. The normalization constant ε_0 is set by the condition

$$\int_{-\pi}^{+\pi} \hat{f}_p(x) dx = 1. \quad (\text{A14})$$

The coefficients $\{a_1, \dots, a_p\}$ are now identified on the basis that the characteristic function of the PDF estimate $\hat{f}_p(x)$ should match the first p terms of the sample characteristic function exactly, namely,

$$\hat{\alpha}_k = \hat{\phi}_k = \int_{-\pi}^{+\pi} \hat{f}_p(x) e^{jkx} dx, \quad 1 \leq k \leq p. \quad (\text{A15})$$

To this purpose, let us compute the quantity $\sum_{k=0}^p a_k \hat{\alpha}_{l-k}$. Recalling that $A(z)$ is analytic in the unit circle and making use of Cauchy's residue theorem, we obtain

$$\sum_{k=0}^p a_k \hat{\phi}_{l-k} = 0, \quad 1 \leq l \leq p, \quad (\text{A16})$$

$$\sum_{k=0}^p a_k \hat{\phi}_{-k} = \varepsilon_0. \quad (\text{A17})$$

Equation (A16) fixes the values of $\{a_1, \dots, a_p\}$ and gives the Yule-Walker equations [Eq. (17)]. The solution is unique provided that

$$\det \begin{bmatrix} \hat{\phi}_0 & \cdots & \hat{\phi}_{-p+1} \\ \vdots & & \vdots \\ \hat{\phi}_{p-1} & \cdots & \hat{\phi}_0 \end{bmatrix} \neq 0. \quad (\text{A18})$$

The latter condition is verified except when a repetitive pattern occurs in the characteristic function. If this happens then the order p should simply be chosen to be less than the periodicity of this pattern.

Besides its positivity, the solution we obtain has a number of useful properties. First, note that all the terms of its characteristic function can be computed recursively by

$$\begin{bmatrix} \hat{\alpha}_{k+1} \\ \hat{\alpha}_k \\ \vdots \\ \hat{\alpha}_{k-p+2} \end{bmatrix} = \begin{bmatrix} -a_1 & -a_2 & \cdots & -a_p \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} \hat{\alpha}_k \\ \hat{\alpha}_{k-1} \\ \vdots \\ \hat{\alpha}_{k-p+1} \end{bmatrix}, \quad (\text{A19})$$

in which the starting condition is set by the p first values of $\hat{\phi}_k$. From this recurrence relation the asymptotic behavior of $\hat{\phi}_k$ as $k \rightarrow \infty$ can be probed by diagonalizing the state space matrix in Eq. (A19). The eigenvalues of this matrix are the roots $\{1/z_1^*, \dots, 1/z_p^*\}$ (called poles), which by construction are all inside the unit disk. Therefore

$$\lim_{k \rightarrow \infty} |\phi_k| \sim e^{\lambda k}, \quad (\text{A20})$$

where λ is related to the largest root and is always negative since

$$\lambda = \max_k \ln \left| \frac{1}{z_k^*} \right| < 0. \quad (\text{A21})$$

This exponential falloff of the characteristic function explains why the resulting PDF is relatively smooth.

Now that we have found a solution in terms of a $[0, p]$ Padé approximant, it is legitimate to ask whether a $[q, p]$ approximant of the type

$$\hat{f}_{q,p}(x) = \frac{|d_0 + d_1 e^{-jx} + \cdots + d_q e^{-jqx}|^2}{|1 + a_1 e^{-jx} + \cdots + a_p e^{-jpx}|^2} \quad (\text{A22})$$

could not bring additional flexibility and hence provide a better estimate of the PDF. Again, we exploit the analogy with spectral density estimation, in which the equivalents of $[q, p]$ Padé approximants are obtained with autoregressive moving average (ARMA) models. The superiority of ARMA over AR models is generally agreed upon [32], although the MISE does not firmly establish it [17]. Meanwhile we note that there does not seem to exist a simple variational principle, similar to that of the likelihood maximization, which naturally leads to a $[q, p]$ Padé approximant of the PDF.

- [1] R. A. Tapia and J. R. Thompson, *Nonparametric Probability Density Estimation* (The Johns Hopkins University Press, Baltimore, 1978).
- [2] B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London, 1986).
- [3] A. Izenman, J. Am. Stat. Assoc. **86**, 205 (1991).
- [4] D. H. Wolpert and D. R. Wolf, Phys. Rev. E **52**, 6841 (1995).
- [5] P. L. G. Ventzek and K. Kitamori, J. Appl. Phys. **75**, 3785 (1994).
- [6] D. L. Donoho, in *Wavelet Analysis and Applications*, edited by Y. Meyer and S. Roques (Éditions Frontières, Gif-sur-Yvette, 1993) pp. 109–128.
- [7] N. N. Cencov, Sov. Math. Dokl. **3**, 1559 (1962).
- [8] S. C. Schwartz, Ann. Math. Stat. **38**, 1262 (1967).
- [9] R. A. Kronmal and M. Tarter, J. Am. Stat. Assoc. **63**, 925 (1968).
- [10] B. R. Crain, Ann. Statist. **2**, 454 (1974).
- [11] A. Pinheiro and B. Vidakovic, Comp. Stat. Data Analysis (to be published).
- [12] U. Frisch, *Turbulence* (Cambridge University Press, Cambridge, England, 1995).
- [13] L. Devroye and L. Györfi, *Nonparametric Density Estimation, the L_1 View* (Wiley, New York, 1985).
- [14] K. Nanbu, Phys. Rev. E **52**, 5832 (1995).
- [15] *Padé Approximants*, edited by P. R. Graves-Morris (The Institute of Physics, London, 1972).
- [16] M. B. Priestley, *Spectral Analysis and Time Series* (Academic Press, London, 1981).
- [17] *Nonlinear Methods for Spectral Analysis*, edited by S. Haykin (Springer-Verlag, Berlin, 1979).
- [18] D. B. Percival and A. T. Walden, *Spectral Analysis for Physical Applications* (Cambridge University Press, Cambridge, England, 1993).
- [19] J.-P. Carmichael, Ph.D. thesis, State University of New York, Buffalo, 1976.
- [20] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing* (Prentice-Hall, Englewood Cliffs, NJ, 1989).
- [21] This is strictly speaking not the characteristic function, but its complex conjugate.
- [22] R. W. Johnson and J. E. Shore, IEEE Trans. Acoust., Speech, Signal Process. **32**, 129 (1984).
- [23] K. N. Berk, Ann. Statist. **2**, 489 (1974).
- [24] E. Parzen, IEEE Trans. Autom. Control. **AC-19**, 723 (1974).
- [25] J. D. Scargle, Astrophys. J. **343**, 874 (1989).
- [26] The width of the interval actually has little impact on the outcome as long as it is much less than 2π . The width and its asymmetry could possibly be tailored to the PDF in order to improve the fit a little.
- [27] The simplest one consists in symmetrizing the PDF: the initial sample $\{x_1, x_2, \dots, x_n\}$ is transformed into a new one which is twice as large $\{x_1, \dots, x_n, 2a-x_1, \dots, 2a-x_n\}$, and covers the interval $[2a-b, b]$. The PDF of this new sample is computed and one-half of it is kept to obtain the desired result. This procedure doubles the volume of data, but the computational cost remains approximately the same since one-half only of the data is actually needed to estimate the characteristic function.
- [28] S. Kullback and R. A. Leibler, Ann. Math. Stat. **22**, 79 (1951).
- [29] C. Beck and F. Schlögl, *Thermodynamics of Chaotic Systems* (Cambridge University Press, Cambridge, England, 1993).
- [30] It can be readily verified that this solution indeed corresponds to a maximum.
- [31] D. E. Smylie, G. K. C. Clarke, and T. J. Ulrich, Methods Comput. Phys. **13**, 391 (1973).
- [32] The main advantage of ARMA models is their ability to model spectral densities that vanish locally.