

Low-Pass Filter Design using Locally Weighted Polynomial Regression and Discrete Prolate Spheroidal Sequences

Tommaso Proietti and Alessandra Luati

Abstract The paper concerns the design of nonparametric low-pass filters that have the property of reproducing a polynomial of a given degree. Two approaches are considered. The first is locally weighted polynomial regression (LWPR), which leads to linear filters depending on three parameters: the bandwidth, the order of the fitting polynomial, and the kernel. We find a remarkable linear (hyperbolic) relationship between the cutoff period (frequency) and the bandwidth, conditional on the choices of the order and the kernel. The second hinges on a generalization of the maximum concentration approach, leading to filters related to discrete prolate spheroidal sequences (DPSS). We propose a new class of low-pass filters that maximize the concentration over a specified frequency range, subject to polynomial reproducing constraints. The design of generalized DPSS filters depends on three parameters: the bandwidth, the polynomial order, and the concentration frequency. We discuss the properties of the corresponding filters in relation to the LWPR filters, and illustrate their use for the design of low-pass filters by investigating how the three parameters are related to the cutoff frequency.

Keywords Nonparametric statistical methods; Kernels; Concentration.

1 Introduction

Trend filters that arise from fitting a locally weighted polynomial to a time series have a well established tradition in time series analysis and signal extraction, see, e.g., Loader (1999). Locally weighted polynomial regression generates finite impulse response filters (LWPR filters, henceforth), that are widely applied in practice.

Tommaso Proietti
University of Rome “Tor Vergata”, S.E.F. e Me.Q., e-mail: tommaso.proietti@uniroma2.it

Alessandra Luati
University of Bologna, Department of Statistics, e-mail: alessandra.luati@unibo.it

LWPR are trend extraction filters, and as such they somehow pass low frequency components and reduce the amplitude of high frequency ones.

The first aim of the paper is to enforce the interpretation of LWPR filters as low-pass filters and to illustrate how the cutoff period (or frequency) is related to three key ingredients: bandwidth, polynomial order and kernel. We define the cutoff period P_c as the smallest periodicity of the fluctuations that are passed i.e. that are preserved to a great extent, by the filter, whereas the fluctuations with smaller period are compressed. Conversely, the cutoff frequency in radians $\omega_c \in (0, \pi)$ is the maximum frequency in radians that is preserved by the filter, whereas higher frequencies are suppressed. The two quantities are inversely related, as $P_c = 2\pi/\omega_c$. An ideal low-pass filter has unit transfer function (TF) in the frequency range $(0, \omega_c)$, and zero TF elsewhere (see Percival and Walden, 1993). For any LWPR filter, we derive the underlying cutoff frequency by least squares filter design principles, as the frequency at which a given LWPR filter provides the best approximation to an ideal low pass filter. Interestingly, we find out that given the kernel and polynomial order, the cutoff period is linearly related to the bandwidth; the intercept and slope of the linear relationship depend on the the order of the polynomial and the kernel. This tight relationship is a mixed blessing: on the one hand, for any LWPR we can determine a reference period or frequency, which would enable us, say, to pick up the correct combination of bandwidth, order and kernel that produces a trend filter passing all the fluctuation with period greater than 8, or 10, years (the deviations of the observations from this component would be an high-pass component that is often identified as the business cycle). On the other hand, the relationship is also a most important limiting factor in the design of a LWPR low-pass filter. For instance, it will turn out that a local linear polynomial fit can produce a low-pass filter with a small cutoff period, such as one year and a half (e.g. 6 quarters), only when it uses a very small bandwidth, which is nevertheless detrimental to the reliability of the estimates (the accuracy of the trend estimator will be poor). If the bandwidth is increased then the cutoff period necessarily increases.

The second contribution of the paper is to propose a generalization of a well-known class of filters that are designed by using the discrete prolate spheroidal sequences (DPSS). The DPSS approach to filter design is based on the maximization of the concentration of a filter at a reference frequency, denoted ϖ . See Tufts and Francis (1970), Slepian (1978) and Lii and Rosenblatt (2008) for recent applications. A DPSS filter reproduces a constant or linear function of time. The filter that solves the traditional concentration problem is given by a zeroth order discrete prolate spheroidal sequence. We generalize the DPSS approach by imposing the constraint that the filter reproduces a polynomial trend of any degree; we thus obtain a class of DPSS filters that depend on three parameters, the bandwidth, the order of the polynomial and the concentration frequency, ϖ . We then discuss the interpretation of the latter. Our interpretation differs from the mainstream one: in fact, we show that the choice of the concentration frequency does not really differ from the choice of the kernel. Finally, we discuss the properties of the corresponding filters in relation to the LWPR filters, and illustrate their use for the design of low-pass filters by investigating how the three parameters are related to the cutoff frequency.

2 LWPR as low-pass filters

Let y_t denote a time series measured at discrete and equally spaced time points. The series can be decomposed as $y_t = \mu_t + \varepsilon_t$, where μ_t is the underlying trend, and ε_t is the noise. We further assume that μ_t is a smooth but unknown deterministic function of time, which can be approximated in a neighborhood of time t by a polynomial of degree p of the time distance, j , between y_t and the neighboring observations y_{t+j} , $j = 0, \pm 1, 2, \dots, h$. This enables to write $\mu_{t+j} \approx m_{t+j}$, with $m_{t+j} = \beta_0 + \beta_1 j + \dots + \beta_p j^p$, $j = 0, \pm 1, \dots, \pm h$. We assume throughout that we are interested in the estimation of the trend at an interior point, and we will not be concerned with the treatment of end points. Provided that $p \leq 2h$, the $p+1$ unknown coefficients β_k , $k = 0, \dots, p$, can be estimated by the method of weighted least squares (WLS), which consists of minimising with respect to the β_k 's the objective function:

$$S(\hat{\beta}_0, \dots, \hat{\beta}_p) = \sum_{j=-h}^h \kappa_j \left(y_{t+j} - \hat{\beta}_0 - \hat{\beta}_1 j - \dots - \hat{\beta}_p j^p \right)^2. \text{ In matrix notation, the local polynomial approximation can be written as } \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}), \text{ where}$$

$\mathbf{y} = [y_{t-h}, \dots, y_t, \dots, y_{t+h}]'$, $\boldsymbol{\varepsilon} = [\varepsilon_{t-h}, \dots, \varepsilon_t, \dots, \varepsilon_{t+h}]'$, the $k+1$ -th column of $\mathbf{X}$ is the $[(-h)^k, \dots, -(h-1)^k, (h-1)^k, (h)^k]'$, $k = 0, \dots, p$, and $\boldsymbol{\beta} = [\beta_0, \beta_1, \dots, \beta_p]'$. Defining $\mathbf{K} = \text{diag}(\kappa_h, \dots, \kappa_1, \kappa_0, \kappa_1, \dots, \kappa_h)$, the WLS estimate of the coefficients is $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{K}\mathbf{X})^{-1}\mathbf{X}'\mathbf{K}\mathbf{y}$. In order to obtain $\hat{m}_t = \hat{\beta}_0$, we need to select the first element of the vector $\hat{\boldsymbol{\beta}}$. Hence, denoting by $\mathbf{e}_1$ the $p+1$ vector $\mathbf{e}_1' = [1, 0, \dots, 0]$, $\hat{m}_t = \mathbf{e}_1' \hat{\boldsymbol{\beta}} = \mathbf{e}_1' (\mathbf{X}'\mathbf{K}\mathbf{X})^{-1} \mathbf{X}'\mathbf{K}\mathbf{y} = \mathbf{w}'\mathbf{y} = \sum_{j=-h}^h w_j y_{t-j}$, which expresses the estimate of the trend as a linear combination of the observations with coefficients

$$\mathbf{w}' = \mathbf{e}_1' (\mathbf{X}'\mathbf{K}\mathbf{X})^{-1} \mathbf{X}'\mathbf{K}. \quad (1)$$

The linear combination yielding our trend estimate is often termed a (linear) *filter*, and the weights w_j constitute its impulse responses. The latter are time invariant and carry essential information on the nature of the estimated signal; they enjoy two important properties, symmetry, $w_{-j} = w_j$, and reproduction of p -th degree polynomials. As far as the second is concerned, from (1) we have that $\mathbf{X}'\mathbf{w} = \mathbf{e}_1$. As a consequence, the filter $\mathbf{w}$ is said to preserve a deterministic polynomial of order p , which means that if the series is $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}$ then the filter will reproduce it exactly, i.e. $\hat{m}_t = \beta_0 = y_t$.

The LWPR filter are low-pass filters, i.e., they leave almost unchanged low frequency components, such as the trend, and attenuate high frequency fluctuations associated with the noise. This is evidenced by the transfer function of the filter

$$G_{h,p,\kappa}(\omega) = \sum_{j=-h}^h w_j e^{-i\omega j}$$

where i is the imaginary unit and $\omega \in (-\pi, \pi)$ is the frequency measured in radians. Due to the symmetry of the weights w_j , $G_{h,p,\kappa}(\omega)$ is real; $|G_{h,p,\kappa}(\omega)|$ is the gain of the filter.

It should be recalled that an ideal low-pass filter retains only the low frequency fluctuations in the series and reduces the amplitude of fluctuations with frequencies higher than a cutoff frequency ω_c . Its transfer function takes the following form: for $\omega \in (0, \pi)$,

$$G_{lp}(\omega) = \begin{cases} 1 & \text{if } \omega \leq \omega_c \\ 0 & \text{otherwise.} \end{cases}$$

As it is well known the ideal filter is available analytically, but unfeasible. Taking the inverse Fourier transform, yields the weights of the ideal filter:

$$w_j = \begin{cases} \frac{\sin(\omega_c j)}{\pi j}, & j = \pm 1, 2, \dots, \\ \frac{\omega_c}{\pi} & j = 0 \end{cases}$$

The filter can be approximated by truncating the weights at lead and lag h , and rescaling them so that their sum is one.

We now derive the cutoff frequency associated with any LWPR filter by a least squares filter design approach. The latter is usually applied for determining the filter which minimizes the squared modulus of the discrepancy between its transfer function, $G_{h,p,\kappa}(\omega)$ and that of the ideal low pass filter, $G_{lp}(\omega)$, for a given cutoff frequency ω_c . See Tufts and Francis (1970) and Percival and Walden (1993, sec. 5.8). In our problem, instead, the filter weights are given and the cutoff is unknown. The latter will be determined as the frequency

$$\omega_c = \operatorname{argmin} \{D(\omega_c; h, p, \kappa)\}, \quad D(\omega_c; h, p, \kappa) = \int_0^\pi |G_{h,p,\kappa}(\omega) - G_{lp}(\omega)|^2 d\omega.$$

for given h, p , and kernel κ . The solution is straightforward: rewriting

$$D(\omega_c; h, p, \kappa) = \int_0^{\omega_c} |G_{h,p,\kappa}(\omega) - 1|^2 d\omega + \int_{\omega_c}^\pi |G_{h,p,\kappa}(\omega)|^2 d\omega,$$

and differentiating with respect to ω_c , the first order condition for the above problem is

$$\sum_{j=-h}^h w_j e^{-i\omega_c j} = \frac{1}{2}, \quad (2)$$

i.e. we need to locate the frequency at which the transfer function is $1/2$. The solution to equation (2) is obtained as the phase of the conjugate pair of roots with unit modulus of the polynomial $\sum_j w_j x^j$.

Figure 1 displays the result for the LWPR filters most commonly in use. The left panel shows the cutoff frequency (vertical axis) corresponding to a particular bandwidth h : as h increases, for given p and kernel function, the cutoff frequency is pushed towards zero, i.e we obtain a filter that produces a smoother trend. It is interesting to present the same results in terms of the cutoff period, $P_c = 2\pi/\omega_c$, since this enables to discover a remarkable linear relationship between cutoff period and bandwidth. For instance, for the Henderson filter, $P_c = 1.4 + 1.1h$, so that when $h = 12$, $P_c = 14.6$. It should be noticed that the intercept and slope depend on the

order of the polynomial and the kernel. If we move from a linear to a cubic fit, *ceteris paribus*, i.e. using the same h and kernel, the cutoff period decreases. Secondly, higher order kernels yield lower cutoff periods for fixed h and p . Figure 1 provides an effective tool for designing the LWPR filter that is appropriate for a particular cutoff frequency or period: drawing an horizontal straight line at a particular value and looking at the intersection with the hyperboles (left panel) or lines (right panel) would make immediately available the bandwidth, order and kernel of the relevant LWPR filter. We provide two applications in the next paragraph.

2.1 Comparison with the Baxter and King filter

Baxter and King (1999, BK henceforth) have popularized the use of the band-pass filter for cycle extraction in economic time series using a least squares approximation to the ideal filter that leads to a finite impulse response filter with a half bandwidth of 3 years, e.g. $h = 12$ when the data are quarterly. A band pass filter is obtained by subtracting two low-pass filters. Band-pass filtering is widely applied various fields, among which economics, to extract signals with given features. For instance, business cycle theory is concerned in the fluctuations with periodicity ranging from 1 year and half to eight years.

When the data are quarterly, the approximation of a low-pass filter with cutoff frequency $\omega_c = 0.2$ ($P_c = 32$, i.e. eight years of quarterly data) as weights that are a scaled version of (2) for $h = 12$. Figure 2 plot this weights along the weights of the local linear LWPR using the Epanechnikov filter and the same bandwidth. The latter




Fig. 1 Relationship between bandwidth (horizontal axis) and cutoff frequency ω_c (left panel), and cutoff period, $P_c = 2\pi/\omega_c$, (right panel) for different LWPR filters.

is characterized by the same cutoff period, and thus the gains of the two filters, plotted in the central panel, are virtually indistinguishable.

Our results on the relationship between LWPR and the cutoff frequency also suggest that the local cubic polynomial regression filters ($h = 24, p = 3$, Epanechnikov) and ($h = 28, p = 3$, Henderson) are characterized by about the same cutoff frequency. The last panel of figure 2 compares their gain with that of the BK approximation to the ideal low-pass filter, showing that the two filters provide a better approximation. However, the higher polynomial order and bandwidths account for the attenuation of both leakage and the Gibbs phenomenon, i.e. the side lobes. If h is increased also for BK, then the filters behave almost identically.

The comparison with the BK filter is also useful to illustrate a possible limitation of LWPR filters, as far as the design of low-pass filters is concerned, which is a consequence of the stringent relationship between bandwidth and cutoff period. As mentioned above, BK construct a cycle filter by subtracting from the low-pass filter with cutoff period $P_c = 32$ quarters the low-pass filter with cutoff period equal to 6 quarters. The latter provides a trend estimate suppressing all the high frequency fluctuations, such as seasonal effect, with periodicity up to one year and half. Both BK low-pass filters have $h = 12$ (three years of quarterly data).

As we have shown, there exists an LWPR filter with $h = 12$ that has cutoff 32 quarters; however, we are unable to find a filter with $\bar{P} = 6$ and $h = 12$. From figure 1 it is clear that when the cutoff period is small, such as $P_c = 6$, the local linear polynomial filters produce a too smooth trend estimate, even for small h . The only LWPR filters that get close to this cutoff are ($h = 4, p = 3$, Henderson) and ($h = 4, p = 3$, Epanechnikov). Figure 3 displays the filter weights and the gains of the two filters and compares them with the BK filter with $\omega_c = \pi/16$ and bandwidth 12 (left panels) and BK with bandwidth 4. Only in the second case we get filters with comparable properties, at the cost of lowering the bandwidth to a very small values.




Fig. 2 Comparison with Baxter and King low-pass filter with $P_c = 32$.

A key consequence is that we cannot design a band-pass filter from two LWPR low-pass filters using the same h , p and kernel. At least one of the elements has to vary. As the bandwidth is the most important factor in determining the cutoff period, h qualifies as the key factor in band-pass filter design.

3 Generalized local polynomial DPSS filter design

We turn our attention to a different approach to the design of a low-pass filter which can be adapted to the problem of fitting a local polynomial. This adaptation configures the main contribution of the paper to this strand of literature.

Let $\mathbf{w}$ denote a filter characterized by the transfer function $G(\omega)$, $\omega \in [0, \pi]$; then its concentration at frequency $\varpi \in (0, \pi)$ is defined as

$$\beta^2(\varpi) = \frac{\int_0^{\varpi} |G(\omega)|^2 d\omega}{\int_0^{\pi} |G(\omega)|^2 d\omega}. \quad (3)$$

The concentration can be expressed as a ratio of quadratic forms,

$$\beta^2(\varpi) = \frac{\mathbf{w}' \mathbf{A}(\varpi) \mathbf{w}}{\mathbf{w}' \mathbf{w}}, \quad (4)$$

where $\mathbf{A}(\varpi)$ is the symmetric and positive definite matrix whose generic ij -th element, $i, j = 1, \dots, 2h + 1$, is




Fig. 3 Comparison with Baxter and King low-pass filter with $P_c = 6$.

$$a_{ij}(\varpi) = \begin{cases} \frac{\sin(\varpi(i-j))}{\pi(i-j)} & \text{for } i \neq j \\ \varpi/\pi & \text{for } i = j \end{cases}$$

as follows by replacing $|G(\omega)|^2 = G(\omega)G^*(\omega) = \sum_{i=-h}^h w_i^2 + 2\sum_{i=-h}^h \sum_{j=-h}^{i-1} w_i w_j \cos((i-j)\omega)$ in (3) and then performing the integrations; $G^*(\omega)$ denotes the complex conjugate of $G(\omega)$.

A strategy for designing filters is, for a given bandwidth h , to choose the filter that maximizes the concentration at a given concentration frequency, under the constraint that the weights sum to one. Under this condition the filter is capable of reproducing a polynomial of order 0, 1. Hence, given ϖ , the filter weights result from the solution of the following constrained maximization problem:

$$\max_{\mathbf{w}_{j,j=-h,\dots,h}} \beta^2(\varpi) \text{ subject to } \mathbf{w}'\mathbf{i} = 1,$$

where $\mathbf{i} = [1, \dots, 1]'$. The quantity (4) is a Rayleigh quotient that reaches its maximum when it is equal to the largest eigenvalue of $\mathbf{A}(\varpi)$, let us call it $\lambda_0(\varpi)$. The corresponding eigenvector, denoted $\mathbf{v}_0$, is the filter that maximizes the concentration (3) or (4) at the concentration frequency ϖ , when no restrictions are imposed to the weights except, possibly, to be of unitary norm. The filter weights arise from the normalization: $\mathbf{w} = \mathbf{v}_0/(\mathbf{v}_0'\mathbf{i})$. The (unit norm) vector $\mathbf{v}_0$ is a subsequence of length $2h+1$ of a zeroth-order discrete prolate spheroidal sequence (DPSS). Its direct computation from the spectral decomposition $\mathbf{A}(\varpi)$ is computationally unattractive and unstable, but it has been shown that it can be derived as the eigenvector associated with the largest eigenvalue of a symmetric tridiagonal matrix, see Slepian (1978).

The DPSS approach described above generates a filter that passes a local constant/linear polynomial, m_{t+j} , for $p = 0, 1$; this is a simple consequence of the constraint $\mathbf{w}'\mathbf{i} = 1$. It seems therefore natural to propose a generalization by introducing further linear constraints on the weights that aim at the reproduction of higher order polynomials. We show below that there is a solution to this problem. The price to be paid is that there is no closed form solution to this problem and the solution has to be found iteratively.

For any bandwidth h , we aim at designing the filter that maximizes the concentration at a given frequency, under the constraint of reproducing a polynomial of degree p . This produces the following constrained maximization problem:

$$\max_{\mathbf{w}_{j,j=-h,\dots,h}} \beta^2(\varpi) \text{ subject to } \mathbf{w}'\mathbf{X} = \mathbf{e}'_1, \quad (5)$$

Obviously, for $\mathbf{X} = \mathbf{i}$, we are back to the standard problem discussed above.

The objective function to be maximized is

$$\phi(\mathbf{w}, \mu) = \frac{\mathbf{w}'\mathbf{A}(\varpi)\mathbf{w}}{\mathbf{w}'\mathbf{w}} - 2(\mathbf{w}'\mathbf{X} - \mathbf{e}'_1)\mu,$$

where $\mathbf{e}'_1 = [1, 0, \dots, 0]$ and μ is a vector of Lagrange multipliers. The first order conditions give $\mathbf{B}(\varpi)\mathbf{w}(\mathbf{w}'\mathbf{w})^{-1} = \mathbf{X}\mu$ and $\mathbf{w}'\mathbf{X} = \mathbf{e}'_1$, where $\mathbf{B}(\varpi) = \mathbf{A}(\varpi) - \beta^2(\varpi)\mathbf{I}$.

If $\mathbf{w}$ is an eigenvector of $\mathbf{A}(\varpi)$, associated with the eigenvalue $\beta^2(\varpi)$, then $\mathbf{B}(\varpi)$ is singular and the first order conditions imply that μ is the null vector, i.e. the problem collapses to the unconstrained maximization discussed above. On the other hand, if $\mathbf{w}$ is not an eigenvector of $\mathbf{A}(\varpi)$, then $\mathbf{B}(\varpi)$ is nonsingular and its inverse $\mathbf{B}(\varpi)^{-1}$ exists. Rearranging, $\mathbf{w}(\mathbf{w}'\mathbf{w})^{-1} = \mathbf{B}(\varpi)^{-1}\mathbf{X}\mu$ and $\mu'\mathbf{X}'\mathbf{B}(\varpi)^{-1}\mathbf{X} = \mathbf{e}'_1(\mathbf{w}'\mathbf{w})^{-1}$, the solution of (5) is the filter

$$\mathbf{w}' = \mathbf{e}'_1(\mathbf{X}'\mathbf{B}(\varpi)^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{B}(\varpi)^{-1}.$$

It is evident that this is an implicit solution, as $\mathbf{B}(\varpi)$ depends on $\mathbf{w}$ through $\beta^2(\varpi)$.

Hence, the solution can be obtained by the following iterative algorithm.

- i. Start from a symmetric LWPR filter $\mathbf{w}_1$ that satisfies the polynomial reproducing constraints.
- ii. For $r = 2, 3, \dots$,
 1. compute the concentration:

$$\beta_r^2(\varpi) = \frac{\mathbf{w}'_r \mathbf{A}(\varpi) \mathbf{w}_1}{\mathbf{w}'_r \mathbf{w}_r};$$

2. construct the matrix $\mathbf{B}_r(\varpi) = \mathbf{A}(\varpi) - \beta_r(\varpi)^2 \mathbf{I}$;
3. update the solution

$$\mathbf{w}'_r = \mathbf{e}'_1(\mathbf{X}'\mathbf{B}_r(\varpi)^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{B}_r(\varpi)^{-1}. \quad (6)$$

- iii. Repeat until convergence, i.e. for some r the discrepancy $\|\mathbf{w}_r - \mathbf{w}_{r-1}\|^2$ is negligible.

We would like to point out the correspondence of (6) with the expression (1) for the LWPR filters. The weights now arise from generalized least square regression using $\mathbf{B}_r(\varpi)^{-1}$ en lieu of the kernel. We shall refer to $\beta_r^2(\varpi)$ as the *restricted concentration* index. It is interpreted as the maximum concentration at the frequency ϖ that a filter reproducing a polynomial of a degree p can achieve.

For the design of a generalized DPSS filter, three parameters are crucial: the bandwidth, the order of the polynomial it can reproduce without bias, p , and the concentration frequency ϖ . Henceforth we shall write $\text{DPSS}(h, p, \varpi)$. Such a filter has associated a value of the concentration index.

The concentration frequency has often been interpreted as a genuine cutoff frequency. We will argue that this interpretation, which has been responsible for the dismissal of this approach for low-pass filter design, is unwarranted. We are going to illustrate that perhaps ϖ is better interpreted as a kernel. Indeed, this interpretation is in line with the view of DPSS filters as *convergence factors*, supported for instance by Percival and Walden (1993, p. 182). Moreover, the comparison of expression (6) with the general expression of an LWPR filter (1) reinforces such interpretation: the only way ϖ affects the shape of the filter is via the matrix $\mathbf{B}(\varpi)$.

Figure 4 displays the value of percentage concentration, $100 \times \beta^2(\varpi)$, as a function of h and either $\bar{P}$ (top left panel) or $\varpi \in (0, \pi)$ (top right panel), for the standard,

i.e. locally constant/linear, DPSS filters. The value of $\beta^2(\varpi)$ is obtained from the first eigenvalue of the matrix $\mathbf{A}(\varpi)$. To enhance the interpretability of the underlying patterns, we present also the contour plots of $100 \times \beta^2(\varpi)$ on the $(h, \bar{P})$ plane (bottom left panel) and (h, ϖ) plane (bottom right panel). For given h , the concentration is a decreasing function of $\bar{P}$ and an increasing function of ϖ . For small periods (high concentration frequencies), and for $\bar{P} < h$, the DPSS filter is 100% concentrated in the desired frequency range. The concentration is poor, instead, when the concentration period (frequency) is high (small) and the bandwidth is small; the latter provides a limiting factor.

The concentration contours enable to assess the trade-off between h and the concentration frequency: to achieve a given concentration level we can decrease h , providing we increase $\bar{P}$ along a straight line, whose intercept and slope depend on the level of $\beta^2(\varpi)$, or equivalently decrease ϖ along an hyperbole. A remarkable feature is indeed that the combinations of $(h, \bar{P})$ giving the same concentration lie along straight lines, whereas the equi-concentration points describe hyperboles on the (h, ϖ_c) plane. The plot also informs us that if $h = 12$ (e.g. three years of quarterly data, so that 25 consecutive observations are used to estimate the trend), we can estimate with unit concentration fluctuations with period up to 22 time units. The concentration for $\bar{P} = 32$ (eight years of quarterly data) is about 94%. To get a concentration equal to 1 for this concentration period a bandwidth $h \geq 18$ would be needed.

The same considerations hold for the local quadratic/cubic DPSS filter given in (6). The major difference lies in the vertical shift of the concentration, which results from the fact that, for given h and either $\bar{P}$ or ϖ , $\beta_r^2(\varpi) \leq \beta^2(\varpi)$. In other words, if we aim at designing a filter with the same concentration, but that passes a cubic,




Fig. 4 Concentration for standard DPSS filters, as a function of the concentration period and frequency.

rather than linear, polynomial, we have to increase either h or $\bar{P}$, or both. The linear and hyperbolic patterns of the contours are present also in this case.

According to the same least squares design principle that was adopted for LWPR filters, the effective cutoff frequency ω_c of any DPSS($h, p, \bar{\omega}$) filter is computed as the frequency at which the transfer function of the filter is equal to one half. Figure 5 presents the results of such computations; in particular, the right panel displays the contours of the cutoff period $P_c = 2\pi/\omega_c$, where ω_c is the frequency at which the gain of the filter equals 1/2, in the standard DPSS case ($p = 0, 1$), as a function of the bandwidth (horizontal axis) and the concentration period (vertical axis). For convenience, we overlay the contours of $100 \times \beta^2(\bar{\omega})$, arising from the same ($h, \bar{P}$) combination (see also figure 4). The plot provides the pairs ($h, \bar{P}$) that produce the same cutoff period. For instance, to have a cutoff of 26 quarters (6 years and a half) we could choose a large $\bar{P}$ and a small bandwidth or a large h and a small $\bar{P}$. The first solution would however provide a poor concentration, e.g. not greater than 20%. The second solution would provide a filter with excellent properties (100% concentration), but more demanding in terms of data points required to estimate the signal (and more difficult to adapt at the extreme of the sample period).

The right panel displays the relationship between the cutoff period P_c and the bandwidth for the case $p = 0, 1$; the different curves correspond to different values of $\bar{P}$; the relationship is nonlinear, but we may appreciate the fact that as $\bar{P}$ increases the curves converge to a straight line that correspond to the Macaulay LWPR filter with the same bandwidth and $p = 0, 1$. In other words, for a fixed bandwidth, if we let $\bar{P}(\bar{\omega})$ increase (decrease), the cutoff period (frequency) increases (decreases) up to an upper limit, which is the value of the cutoff period for the local linear Macaulay filter.




Fig. 5 DPSS($h, p, \bar{\omega}$) filters. Contour plots of the cutoff period $P_c = 2\pi/\omega_c$ as a function of the bandwidth h , and the concentration period (left panel, case $p = 0, 1$). Cutoff period as a function of bandwidth for different values of $\bar{P}$ (right panel).

This finding actually generalizes to any order p . In general, for a given h the limiting DPSS(h, p, ϖ) filter as ϖ decreases, is the LWPR(h, p , Macaulay) filter. This is illustrated in the same plot by the dashed line which refers to the local cubic Macaulay filter (see also figure 1) and is already coincident with the bandwidth-cutoff period relationship for the DPSS($h, 3, \pi/60$) filter. Another conclusion that we draw from the graph is that the generalized DPSS approach can reveal itself more flexible for the design of band-pass filters, with respect to the LWPR approach. In particular the concentration frequency, or period, can be seen as a device that is responsible for the continuous adaptation of the kernel of the local polynomial fit.

4 Concluding remarks

The paper has first aimed at characterizing local polynomial filters as low-pass filters, by investigating the relationship between the cutoff frequency and the three ingredients of the design: the bandwidth, the order of the polynomial, and the kernel. The conclusions of this investigation provide guidance over the choice of the three parameters for designing a low-pass filter with a particular cutoff frequency. They also point out that the strategy of designing a band-pass filter (e.g. for isolating the business cycle fluctuations) by subtracting two low-pass filters using the same bandwidth, kernel and polynomial order is unfeasible, as at least one of the elements has to change.

Secondly, the paper has proposed a generalized DPSS filter design approach that depends on three parameters: the bandwidth, the order of the polynomial that is reproduced by the filter, and the concentration frequency. The paper has investigated how this three parameters can be combined so as to design a low-pass filter with a preassigned cutoff frequency. The overall conclusion is that the DPSS filters have more flexibility and can provide a more coherent design for band-pass filtering. The preferred strategy would be to contrast two DPSS low-pass filters using the same bandwidth and polynomial order, but different concentration frequencies.

References

1. Baxter, M., King, R.G. (1999). Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series. *The Review of Economics and Statistics*, 81, 575-593.
2. Lii K.S., Rosenblatt M. (2008), Prolate Spheroidal Spectral Estimates, *Statistics and Probability Letters*, 78, 11, 1339-1348.
3. Loader, C. (1999): *SLocal Regression and Likelihood*. Springer-Verlag, New York.
4. Percival, D. and Walden, A. (1993): *Spectral Analysis for Physical Applications*. Cambridge University Press, Cambridge.
5. Slepian, D. (1973): Prolate Spheroidal Wave Functions, Fourier Analysis and Uncertainty - V: The Discrete Case. *The Bell System Technical Journal* 57, 1371-1430.
6. Tufts D.W., Francis J.T. (1970), Designing Digital Low-Pass Filters- Comparison of Some Methods and Criteria, *IEEE Transactions on Audio and Electroacoustics*, 18, 4, 487-494.