Efficient Digital Pre-Filtering For Least-Squares Linear Approximation

Marco Dalai, Riccardo Leonardi and Pierangelo Migliorati

University of Brescia, Via Branze 38, 25123, Brescia (Italy)
{*name.surname*}@ing.unibs.it

Abstract. In this paper we propose a very simple FIR pre-filter based method for near optimal least-squares linear approximation of discrete time signals. A digital pre-processing filter, which we demonstrate to be near-optimal, is applied to the signal before performing the usual linear interpolation. This leads to a non interpolating reconstruction of the signal, with good reconstruction quality and very limited computational cost. The basic formalism adopted to design the pre-filter has been derived from the framework introduced by Blu et Unser in [1]. To demonstrate the usability and the effectiveness of the approach, the proposed method has been applied to the problem of natural image resampling, which is typically applied when the image undergoes successive rotations. The performance obtained are very interesting, and the required computational effort is extremely low.

1 Introduction

Linear interpolation is one of the simplest methods for digital to analog conversion, and it is still applied in applications where the conversion computational cost must be kept under control. Nevertheless, linear interpolation introduces some artifacts, e.g., blurring in image processing, and it is often necessary to use higher degree polynomials to achieve an acceptable level of approximation.

An interesting point is to establish if a signal needs to be actually interpolated or whether a non interpolating least-square approximation could be preferable. In literature the choice seems to have been interpolation, even when it does not represent a requirement and it clearly generates analog reconstructed signals with a much higher mean square error. A significant example is represented by the problem of image resampling when the new pixels are uniformly distributed with respect to the old ones, as in the case of image rotations by angles which are not multiples of $\pi/2$.

The concept of least-squares reconstruction of signals has been well studied in modern sampling theory (e.g., see [2]), and the solution is known to be given by an analog pre-filtering of the signal before sampling. The optimal filter inpulse response depends on the generating function used in the reconstruction step (and is called *dual* of the generating function itself).

In this paper we propose the introduction of a very simple digital filter to be applied to the signal before performing the classical linear interpolation, as an

equivalent of the analog analysis filter which is missing in our implementation. The approach is a direct extension of the results described by Blu and Unser in [1]. We show here that a very simple FIR filter leads to almost the same performance with respect to an optimal analog analysis pre-filter under the hypothesis that the signal has been sampled according to Shannon's sampling theorem. In this paper we will use a normalized frequency representation, so that this hypothesis is equivalent to consider the signal to be bandlimited to frequencies $|f| < 1/2$.

The rest of the paper is organized as follows. Section 2 describes some known results of interpolation and sampling theory, whereas the proposed digital pre-filter is introduced and discussed in Section 3. Experimental results showing the effectiveness of the proposed approach are reported in Section 4. Finally, concluding remarks are drawn in Section 5.

2 Interpolation and Sampling Theory

Given a set of values $s(k)$, corresponding to samples of a signal s at integer points, a linear interpolation constructs a piecewise linear approximation $\tilde{s}$ given, for $k \in N$ and $\delta \in [0, 1[$, by

$$\tilde{s}(k + \delta) = s(k)(1 - \delta) + s(k + 1)\delta$$

This is known to be equivalent, for $t \in \mathbb{R}$, to

$$\tilde{s}(t) = \sum_{k \in \mathbb{Z}} s(k)\, \Lambda(t - k) \tag{1}$$

where $\Lambda(t)$ is the unity triangular pulse defined by $\Lambda(t) = \max(0, 1 - |t|)$.

Equation (1) is a special case of a general model used in modern sampling theory: given a reconstruction function φ a signal approximation can be obtained using the linear expansion

$$\tilde{s}(t) = \sum_{k \in \mathbb{Z}} c_k\, \varphi(t - k) \tag{2}$$

where the coefficients c_k have to be set depending on φ, s, and the desired characteristics of the approximation result. A very useful choice is to construct an *analysis* function $\tilde{\varphi}$, dependent on φ, and consider coefficients c_k of the form

$$c_k = \langle s, \tilde{\varphi}_k \rangle$$

where $\tilde{\varphi}_k(t) = \tilde{\varphi}(t - k)$. This assumption corresponds to filtering s with a filter with impulse response $h(t) = \tilde{\varphi}(-t)$, followed by an ideal sampler to obtain the sequence c_k.

With this formalism in mind, given a certain reconstruction function φ, the approximation $\tilde{s}$ is simply determined by the choice of $\tilde{\varphi}$. If we want to obtain

the least-square approximation solution (in the form of eq. (2)), the analysis function is given by (see [2]) [1]

$$\hat{\tilde{\varphi}} = \frac{\hat{\varphi}(f)}{\hat{a}_\varphi(f)}$$

where $a_\varphi(t)$ is the sampled autocorrelation of the function φ, and thus satisfies

$$\hat{a}_\varphi(f) = \sum_{k \in Z} |\hat{\varphi}(f - k)|^2$$

In this case $\tilde{\varphi}$ is called *dual* of φ and will be denoted with $\overset{\circ}{\varphi}$.

If we are interested to linear approximation, the expression of $\overset{\circ}{\varphi}(f)$ can be easily derived, given that $\varphi(t) = \Lambda(t)$ and $\hat{\varphi}(f) = \text{sinc}^2(f)$. In this case one has

$$\hat{a}_\Lambda(f) = \frac{2}{3} + \frac{1}{3}\cos(2\pi f) \tag{3}$$

and, thus,

$$\hat{\overset{\circ}{\Lambda}} = \frac{3\,\text{sinc}(f)^2}{2 + \cos(2\pi f)}$$

Consider that this function is even and real valued, then $\Lambda(t)$ must be even as well. Thus we have $h(t) = \overset{\circ}{\Lambda}(-t) = \overset{\circ}{\Lambda}(t)$, and we will refer to the dual function $\overset{\circ}{\Lambda}$ as the analysis filter.

3   The Proposed Digital Pre-Filter

Once we have the expression of the analog analysis filter for the least-squares approximation this should be applied to the signal before sampling, in view of a future piecewise linear reconstruction (which is optimal according to the least squares criterion).
Obviously, this cannot be implemented in practice as the signal is normally available already in a digital form. So, the overall processing must be implemented in an equivalent way in the digital domain. Assuming that the sampling rate satisfies Shannon's sampling theorem, a near equivalent FIR pre-filter can be derived and applied to the digital signal providing nearly the same performance as its optimal analog counterpart.
In this way, it will be possible to reconstruct a piecewise linear least-square approximation of the original signal in a very efficient manner.

---

[1] $\hat{\varphi}$ represents the Fourier transform of the function φ. All functions with the superscript ^ will in what follows always identify the Fourier transform of its argument, unless it is explicitely stated differently in the text

3.1 Filter

We choose to limit the number of taps of the near equivalent FIR filter to 5. This could be considered a completely arbitrary choice, but it will be shown that the quality of the approximation is practically unaffected by such limitation. In addition, this allows to maintain a very low computational cost. Longer filters have been considered without obtaining a substantial reduction in the mean square error.

As previously mentioned, the FIR digital filter should behave similarly to its Analog counterpart $\overset{\circ}{\Lambda}$. The design approach basically equates the first terms of the Taylor series expansion of both frequency representation of the two filters around $f = 0$. Given the symmetries of $\overset{\circ}{\Lambda}$, $h[n]$ has to be even shaped and satisfy a unity gain at $f = 0$. Thus its generic expression in the Z transform domain is given by

$$h(z) = 1 - a - b + \frac{a}{2}(z^1 + z^{-1}) + \frac{b}{2}(z^2 + z^{-2})$$

The associated Discrete Time Fourier Transform is thus

$$\hat{h}(f) = 1 - a - b + a\cos(2\pi f) + b\cos(4\pi f)$$

Using now its Taylor series expansion and limiting it to the fourth order approximation, one can rewrite

$$\hat{h}(f) = 1 + (2a - 8b)\pi^2 f^2 + \frac{32b - 2a}{3}\pi^4 f^4 + o(f^5)$$

Similarly the Taylor Series representation of the dual function is given by

$$\overset{\circ}{\hat{\Lambda}}(f) = 1 + \frac{1}{3}\pi^2 f^2 + \frac{2}{45}\pi^4 f^4 + o(f^5)$$

Equating the coefficients of f^2 and f^4 we obtain a linear system with two unknown, the solution of which leads to $a = -11/45$ and $b = 7/360$. Thus the filter transfer function is given by

$$h(z) = \frac{49}{40} - \frac{11}{90}(z^1 + z^{-1}) + \frac{7}{720}(z^2 + z^{-2})$$

The joint concatenation of this digital filter with a linear interpolator can also be viewed as an approximation process in the form of equation (2). The coefficients c_k correspond to the samples $s(k)$, and the reconstruction function φ is given by a linear combination of $\Lambda(t)$, $\Lambda(t \pm 1)$ and $\Lambda(t \pm 2)$ leading to the impulse response shown in fig 1.

3.2 Approximation error

Equipped with the analytical expression of our filter, it is important to compare its performance with that obtained by the optimal solution. In [3] and [1], the
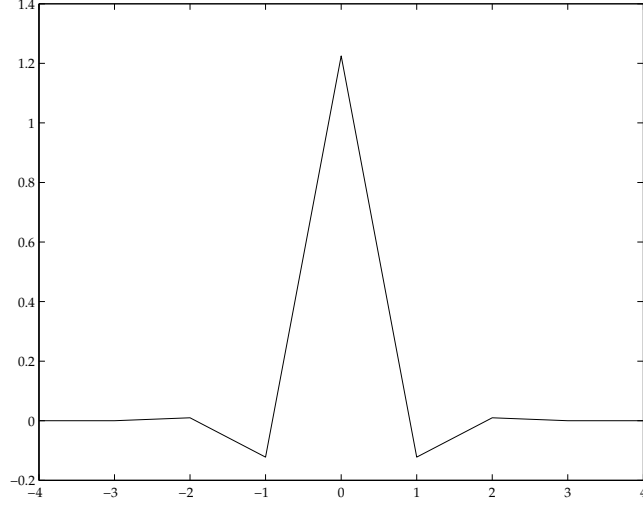
Fig. 1. Equivalent reconstruction function of the FIR pre-filtering method

authors gave extremely useful mathematical tools for the study of the mean square error behavior in approximation techniques. In particular it is shown that, if T is the sampling rate, a very good estimation of the mean square approximation error is given by

$$\eta_s(T) = \left(\int_{-\infty}^{\infty} |\hat{s}(f)|^2 E_{\tilde{\varphi},\varphi}(fT) df \right)^{1/2}$$

where the $E_{\tilde{\varphi},\varphi}(f)$ is called the *error kernel* function and is given by

$$E_{\tilde{\varphi},\varphi}(f) = |1 - \hat{\tilde{\varphi}}(f)^* \hat{\varphi}(f)|^2 + |\hat{\tilde{\varphi}}(f)|^2 \sum_{n \neq 0} |\hat{\varphi}(f+n)|^2 \qquad (4)$$

or equivalently

$$E_{\tilde{\varphi},\varphi}(f) = 1 - \frac{|\hat{\varphi}(f)|^2}{\hat{a}_\varphi(f)} + \hat{a}_\varphi(f)|\hat{\tilde{\varphi}}(f) - \hat{\overset{\circ}{\varphi}}(f)|^2 \qquad (5)$$

For our purpose we set T to 1, as we are using a normalized frequency representation. We have therefore to compare the error kernel for the least square approximation using piecewise linear functions, the classic linear interpolation, and our near optimal solution.

For the first case, we have that $\tilde{\varphi} = \overset{\circ}{\Lambda}$ and, calling $E_{\min} = E_{\overset{\circ}{\Lambda},\Lambda}$, from (5) and (3),
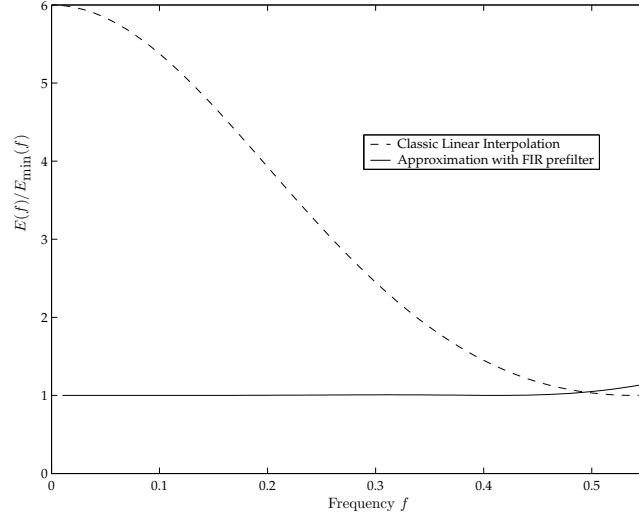
$$E_{\min}(f) = 1 - \frac{3\operatorname{sinc}(f)^4}{2 + \cos(2\pi f)}$$

For the classic linear interpolation, instead, $\tilde{\varphi}(t) = \delta(t)$ and thus, from (4),

$$E_{\delta,\Lambda}(f) = \frac{5}{3} + \frac{1}{3}\cos(2\pi f) - 2\operatorname{sinc}^2(f)$$

Finally, for the proposed FIR prefilter, $\tilde{\varphi} = h$ and, from (4), we obtain

$$E_{h,\Lambda}(f) = 1 + h^2(f)\left(\frac{2}{3} + \frac{1}{3}\cos(2\pi f)\right) - 2h(f)\operatorname{sinc}^2(f)$$
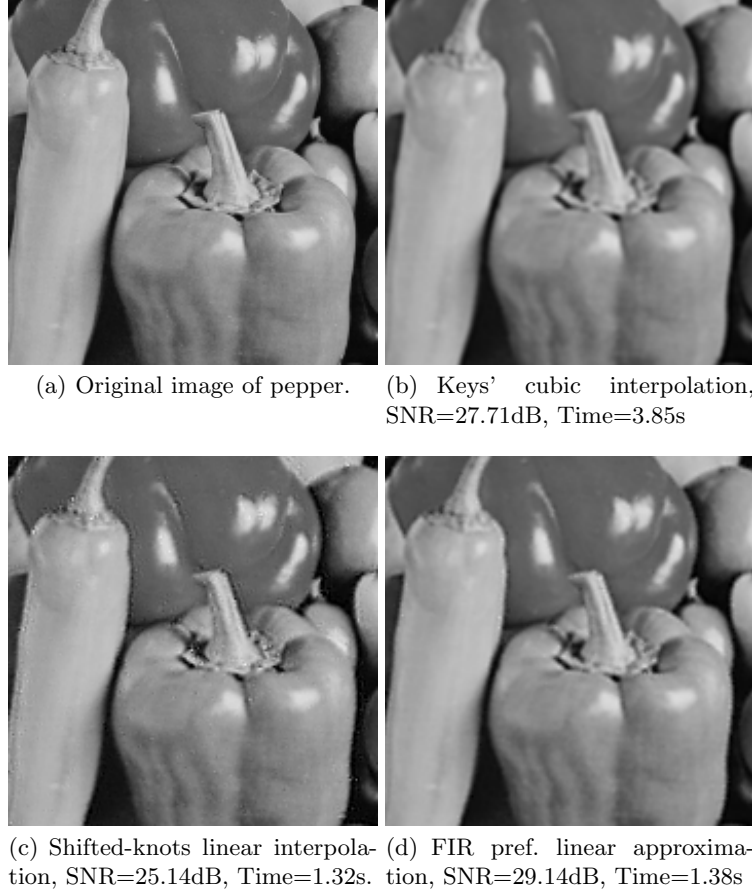

Fig. 2. Error kernels comparison.

Comparing the power series expansion of these three functions, one may verify that $E_{\delta,\Lambda}(f) - E_{\min}(f) = O(f^4)$ while $E_{h,\Lambda}(f) - E_{\min}(f) = O(f^{12})$ which means that for low frequencies the FIR pre-filter leads to much smaller errors with respect to the classical interpolation. In Fig. 2 the normalized error kernels $E(f)/E_{\min}(f)$ are plotted for both methods. It clearly indicates that the proposed FIR pre-filtering method is practically identical to that of the least square solution up to the Nyquist frequency $f = 1/2$. Therefore, if the signal can be considered bandlimited to $|f| < 1/2$ (i.e., Shannon's sampling theorem is satisfied), the solution obtained by FIR pre-filtering can be considered equivalent to the optimal analog solution.

4 Simulation Results

In order to demonstrate the efficiency of our approximation scheme, we compare it with two meaningful interpolation methods, namely Keys' cubic interpolation,

which is the reference for high-quality image reconstruction methods, and the more recent interesting shifted-knots linear interpolation proposed by Blu *et al.* in [4]. We perform an image resampling test, based on successive rotation of 256 gray level images, that is a variant of that used in [4].
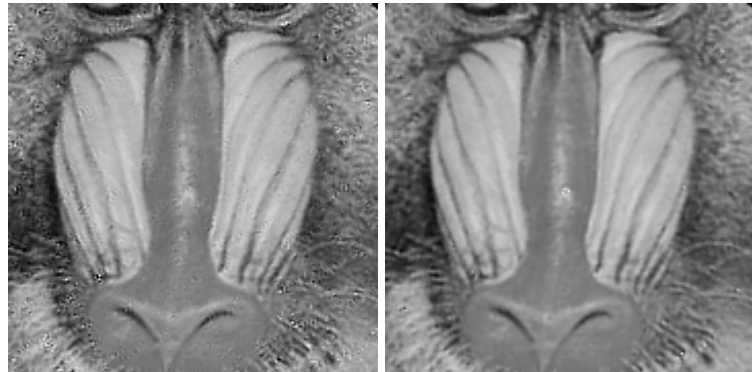

(a) Original image of pepper.

(b) Keys' cubic interpolation, SNR=27.71dB, Time=3.85s

(c) Shifted-knots linear interpolation, SNR=25.14dB, Time=1.32s.

(d) FIR pref. linear approximation, SNR=29.14dB, Time=1.38s

Fig. 3. Results of simulations on the image Pepper.

We randomly choose 14 angles θ_i with $\theta_i \in [-\pi/2, \pi/2]$ for $1 \le i \le 14$ and we set $\theta_{15} = 0$. The value θ_i represents the angular position (with respect to the original one) of the image at the i-th step. The image at step $i + 1$ is obtained by rotating the image at step i by an angle $\theta_{i+1} - \theta_i$. It is clear that at the 15-th step the image is exactly in the initial position and we can compare the rotated image with respect to the original one to evaluate the error introduced by successive resamplings of the rotated rectangular image. All computations have been performed with floating point arithmetic so as to avoid the effects

(a) Original image of pepper.

(b) Keys' cubic interpolation, SNR=20.50dB, Time=3.85s

(c) Shifted-knots linear interpolation, SNR=18.02dB, Time=1.32s

(d) FIR pref. linear approximation, SNR=21.42dB, Time=1.38s

Fig. 4. Results of simulations on the image Pepper.

of quantization. Furthermore, the error has been computed only on the central part of the image so that boundaries effect have been discarded.

This simulation experiment is different from that proposed in [4] for which the image was rotated 15 times by an angle $2\pi/15$ in the same direction. Such a choice, in fact, causes a compensation of the phase distortion potentially introduced (and it is the case for the shifted-knots method) by the approximation system, due to the presence of near-opposite angles during the simulation. This is the reason for which, here, shifted-knots linear interpolation gives smaller SNR values than Keys' cubic one (even if the images can be considered visually more satisfactory), contrarily to what reported in [4].

Figures 3 and 4 show the results of the test for the Pepper and Baboon images, respectively. It is important to note that the FIR-prefilter linear approximation can give higher quality (both visual and with respect to the SNR value) than Keys' cubic interpolation in 1/3 of the computation time. Furthermore, it is interesting to compare the performance of our method to that of the shifted-knots one; the phase distortion of the latter, in fact, leads to a lower SNR value in our rotation test. Visually, the obtained images are not blurred (as with classic linear interpolation) but they are affected by the presence of shot-noise. The least-squares linear approximation, on the contrary, does not introduce this artifact and gives images with very little blur, leading to higher quality for about the same computational cost.

5 Conclusion

In this paper we have presented the idea of using a very simple linear phase FIR pre-filter to compute a least square linear approximation of a digital signal for bandlimited analog signals reconstruction. The digital pre-processing filter is applied to the signal before performing the classic linear interpolation, so as to obtain a non interpolating analog approximation. The usefulness of the approach has been demonstrated by applying it to the problem of resampling rotated images. In this context, the need for an exact interpolation of the original pixels is not a requirement, and the quality of the rotated signal remains very close to the original.

References

1. Blu, T., Unser, M.: Quantitative Fourier analysis of approximation techniques: part I–interpolators and projectors. IEEE Trans. Signal Process. **47**(10) (1999) 2783–2795
2. Unser, M.: Sampling-50 years after Shannon. Proc. IEEE **88**(4) (2000) 569–587
3. Blu, T., Unser, M.: Approximation error for quasi-interpolators and (multi)-wavelet expansions. Appl. Comput. Harmon. Anal. **6**(2) (1999) 219–251
4. Blu, T., Thvenaz, P., Unser, M.: How a simple shift can significantly improve the performance of linear interpolation. Proc. IEEE Int'l Conf. on Image Proc. (2000) III.377–III.380