Signal Processing 114 (2015) 251–264

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Multidimensional digital smoothing filters for target detection

Hugh L. Kennedy

Defence and Systems Institute, University of South Australia, Australia

A R T I C L E I N F O

Article history:
Received 8 October 2014
Received in revised form
11 February 2015
Accepted 8 March 2015
Available online 18 March 2015

Keywords:
Digital filters
Discrete Laguerre Transform
IIR filters
Image analysis
Multidimensional signal processing

A B S T R A C T

Recursive, causal and non-causal, multidimensional digital filters, with infinite impulse responses and maximally flat magnitude and delay responses in the low-frequency region, are designed to negate correlated clutter and interference in the 'background' and to accumulate power due to dim targets in the 'foreground' of a surveillance sensor. Expressions relating mean impulse-response duration, frequency selectivity and group delay, to low-order linear-difference-equation coefficients are derived using discrete Laguerre polynomials and discounted least-squares regression, then verified through simulation.

1. Introduction

Low-pass digital filters, such as those proposed by Savitzky and Golay, with a 'maximally flat' magnitude and delay response, have smoothing properties in the time domain. This 'duality' makes it possible to derive their filter coefficients in either the frequency *or* time domains [1–4]. For instance, Savitzky–Golay filters [5], may be derived by least-squares fitting a polynomial (of degree B) to a sampled input sequence over a finite sliding window to yield low-pass filters with a finite impulse response (FIR). The fitted polynomial resulting from this *analysis* process, is evaluated at the center of the odd analysis window to yield linear-phase (smoothing) filters; evaluation between samples yields fractional-delay (interpolating) filters, evaluation at more recent non-central samples yields filters with a reduced group delay; whereas evaluation at future samples yields predictive (extrapolating) filters. The offset (q) chosen for the evaluation, or *synthesis*, therefore determines the phase response of the filter [6]. Savitzky–Golay differentiators [7], are obtained by differentiating the fitted polynomial prior to evaluation. FIR Savitzky–Golay filters are realized using either non-recursive or recursive structures; however care is required in the latter case to avoid rounding error accumulation due to pole-zero cancellation on the unit circle [6].

'Fading-memory' variants of these 'finite-memory' Savitzky–Golay filters may similarly be derived by performing a least-squares fit with an exponentially-decaying error-weighting function (whose $\mathcal{Z}$ transform has a pole at $\mathcal{Z} = p$, where $p = e^{\sigma}$), yielding recursive structures with an infinite impulse response (IIR) and with stability guaranteed (for all q, if $|p| < 1$) [8–11]. They are commonly used in target tracking systems to overcome problems of divergence experienced by Kalman filters in the presence of model mismatch [9]; however in this context, they are usually restricted to applications where the time interval between target detections is constant and where data association is unambiguous; furthermore, startup transients must be handled properly. These restrictions have recently been addressed in [12]; with an expanding-memory filter used during track establishment and a fading-memory filter used thereafter; measurements are probabilistically weighted, however the revisit interval is assumed constant.

For satisfactory tracking performance at a reasonable computational cost, these types of 'detect-before-track' systems require the signal-to-noise ratio (SNR) to be relatively large so that (1) the density of false-detections – due to clutter, interference and noise – is low, (2) the probability of

---

E-mail address: hugh.kennedy@unisa.edu.au

http://dx.doi.org/10.1016/j.sigpro.2015.03.005
0165-1684/

target detection is high, and (3) the measurement error is low. When these conditions cannot be met, methods that exploit spatiotemporal energy distributions in the underlying sensor 'image', i.e. 'track-before-detect' methods are more appropriate [13–17]. Target confirmation decisions in both detect-before-track and track-before-detect frameworks are usually framed as hypothesis tests, based on a test statistic involving a likelihood ratio. The likelihood functions for the true and false detections generally have simple idealized forms – Gaussian for the target and Poisson for the clutter/noise, in spatial coordinates. The Rician distribution is ideal for modeling intensity distributions because it results in Rayleigh and Gaussian distributions at the low- and high-SNR limits, respectively [14]. When the assumed forms or the estimated parameters of the underlying distributions are inappropriate or in error, severely degraded target detection and tracking performance results [16].

If not handled properly, structured backgrounds have the potential to 'wreak havoc' in detect-before-track and track-before-detect systems alike. While often ignored in theoretical works, correlated clutter/noise is commonplace in long-range surveillance systems involving infrared search and track cameras, high-frequency radar, passive (bi-static) radar and sonar [16], especially when very weak targets are sought.

The use of fading-memory target-tracking filters, derived using discounted least-squares and orthogonal polynomials, has been thoroughly explored in the literature, in a *detect-before-track* context [8–12]; indeed, it is interesting and illuminating that the same filter coefficients are reached from such a variety of different starting points. However, these filters are arguably a more natural solution to the target-tracking problem in a *track-before-detect* context, where uniform sampling rates are guaranteed in the spatial domain (i.e. within a frame or dwell) and expected in the temporal domain, under normal operating conditions (although there may be some 'jitter'). Furthermore, the complication of data association is avoided and special logic need not be implemented to handle initialization and start-up transients, although this may be required in the spatial dimensions if a reduction in the sensor's field of regard is unacceptable. In this paper, multidimensional forms of the filters are used to perform both background whitening *and* target enhancement functions.

Strictly speaking, the proposed algorithm is neither a detect-before-track nor a track-before-detect approach because no attempt is made to establish, and maintain continuity, of target identity. As a consequence, data association is avoided, thus the computational load is constant and data independent, i.e. it does not depend on the density or intensity of the target or clutter; furthermore, the filters are amenable to parallelization because the same operations are applied to every 'cell'. The resulting SNR enhancement should improve the performance and simplify the structure of any 'downstream' detect-before-track stage that follows. The proposed algorithm might therefore be regarded as an 'enhance-before-detect' approach. Non-causal (forward/backward) IIR filters are used in the spatial dimensions, whereas an IIR filter with a tunable group delay, is used in the temporal dimension. The simple premise underlying the derivation of the filters allows them to be intuitively adapted and tuned for a wide range of functions.

Analog filter prototypes are used to design the multidimensional IIR filters in [18,19]; whereas, a direct digital design approach is adopted here. Classical analysis offers the designer an array of well-established relationships to build analog filters; however they do not transfer exactly into the digital domain so a 'sympathetic' discretization method must be chosen to ensure that the intent of the original design is preserved, which adds an extra layer of complexity to the design process.

Matters relating to the use of multidimensional IIR filters, with maximally-flat responses, in enhance-before-detect algorithms, are addressed in this paper: (1) closed-form expressions for the coefficients of low-order linear-difference-equations in terms of the forgetting factor ($\sigma, p = e^{\sigma}$) and the synthesis offset (q) are derived; (2) relationships between these design parameters and the frequency response (magnitude and phase) of the filter are described; (3) ways in which the filter response influences the performance of the enhance-before-detect algorithm are discussed; and (4) a technique for estimating point-target velocity by exploiting the local 'Laguerre spectrum' is proposed. A particular filter arrangement that is very well-suited to enhance-before-detect roles is also presented – a background subtraction stage is cascaded with a foreground accumulation stage; both stages use non-causal filters in the spatial dimensions and causal filters in the temporal dimension.

Not all of the relationships required for the task at hand have been tabulated in the literature, for instance, phase control is omitted in [9], only causal filters are considered in [8], and the discussion in [6] is limited to (recursive) FIR filters, with pole-zero cancellation on the unit circle, which is good for efficiency but bad for immunity to rounding error accumulation. Non-causal IIR smoothers and differentiators are presented in [20] however the treatment is restricted to first- and second-order filters. Frequency-domain properties are not analyzed in [8–11] and the usefulness of analysis-only operations, to yield the Laguerre spectrum [8], is typically overlooked in the modern literature. Recursive analysis-only filters are also derived and applied in this paper.

There are a number of other non-iterative closed-form techniques for deriving the coefficients of low-pass digital filters with maximally-flat responses, that resemble the much-loved monotonic responses of classical, Bessel and Butterworth, analog-filters [21–27].

In Hermann's early treatment of the problem, exactly linear-phase FIR solutions that satisfy magnitude flatness constraints up to a specified derivative order, at frequencies $\omega = 0$ and $\omega = \pm\pi$, are derived [21]. Low-latency low-pass FIR filters may also be designed to satisfy magnitude and group delay flatness constraints at $\omega = 0$ only, which allows the group delay to be varied [22,23]. Using the closed-form expressions in [24], specification of the filter order, the desired group delay, and the number of zeros at $\mathcal{Z} = -1$, yields filters with good phase linearity at low-frequencies and very good high-frequency attenuation; however, the ability to control near-DC gain, i.e. bandwidth and roll-off, is limited for low-order filters. In an extension of this work, 'partially flat' FIR filters with derivative-constraints are investigated in [25].