Recursive Gaussian Derivative Filters

Lucas J. van Vliet, Ian T. Young and Piet W. Verbeek

Pattern Recognition Group of the Faculty of Applied Physics
Delft University of Technology
Lorentzweg 1, 2628CJ Delft, The Netherlands

email: L.J.vanVliet@ph.tn.tudelft.nl
http://www.ph.tn.tudelft.nl

Abstract

We propose a new strategy to design recursive implementations of the Gaussian filter and Gaussian regularized derivative filters. Each recursive filter consists of a cascade of two stable N^{th}-order subsystems (causal and anti-causal). The computational complexity is 2N multiplications per pixel per dimension independent of the size (σ) of the Gaussian kernel. The filter coefficients have a closed-form solution as a function of scale (σ) and recursion order N (N=3,4,5). The recursive filters yield a high accuracy and excellent isotropy in n-D space.

1. Introduction

Gaussian (derivative) filters are used in a wide variety of computer vision tasks. The Gaussian filter is frequently used as a low-pass filter for noise suppression or scale-space construction [1, 2]. Optimal edge detection uses Gaussian regularized derivatives to detect and localize 1-D noisy step edges [3]. Accurate localization of curved edges uses a nonlinear combination of Gaussian derivatives [4]. Geometric object features such as perimeter length in 2-D, curve length and surface area in 3-D, and shape descriptors such as bending energy and Euler number can be measured using Gaussian derivative filters [5]. Local anisotropy and orientation can be measured in a Gaussian smoothed tensor image constructed from the dyadic product of a gradient vector with itself [6, 7].

These applications will benefit from a recursive implementation of Gaussian derivative filters. Recursive filters are fast and scale-independent. Deriche [11-13] has proposed a stable recursive implementation by decomposing the transfer function into a sum (*parallel interconnection*) of a stable causal and a stable anti-causal filter. Each sub-filter of order N requires $2N$ multiplications, i.e. N recursive and N non-recursive. The work presented here is an extension of our previous work [9]. In [9] we derived a stable filter by decomposing the transfer function in a product (*series interconnection*) of a stable causal and a stable anti-causal filter. Each sub-filter of order N required only N multiplications, i.e., N recursive and zero non-recursive.

In this work we use the efficient series (cascade) implementation from [9] which is faster than any other alternative [9]. The major drawbacks of our earlier filter have been addressed by a new design strategy. The improvements are:
- improved accuracy (significantly smaller error for small Gaussians),
- arbitrary order of recursive filter minimizing the L^2 or L^∞ norm,
- closed-form solution of the filter coefficients as a function of σ and N,
- improved isotropy and accuracy for Gaussian derivative filters in *n*-D space.

2. Gaussian filter

The multi-dimensional Gaussian filter can be separated into one-dimensional Gaussians along its main axes. For simplicity we start our analysis with a 1-D Gaussian. The Gaussian filter $g(t;\sigma)$ has a Fourier transform $G(\omega;\sigma)$ which is also a Gaussian

$$g(t;\sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-t^2}{2\sigma^2}} \quad \xleftrightarrow{\mathcal{F}} \quad G(\omega;\sigma) = e^{\frac{-\omega^2\sigma^2}{2}} \quad (1)$$

and where σ is referred to as both the *standard deviation* and the *scale*. The Gaussian filter is the filter which minimizes the product of spatial support and frequency bandwidth [8]. It allows simultaneous localization in the time (space) and (spatial) frequency domain.

3. Discrete-time recursive filter

The goal is to design a discrete recursive Gaussian filter that requires a minimum number of multiplications. A discrete time (space) system is described by its *z*-transform. The Gaussian transfer function is real and even. A real and even transfer function requires that $H^*(z)=H(z)$ and $H(z)=H(z^{-1})$. This is guaranteed by our choice

$$H(z) = K(z) \cdot K^*(z) \cdot K\left(z^{-1}\right) \cdot K^*\left(z^{-1}\right) \qquad (2)$$

As a consequence, all complex-valued poles (and zeroes) appear in quadruples $\{d_i,\ d_i^{-1}, d_i^*, d_i^{*-1}\}$. For real-valued poles $H(z) = K(z)\,K(z^{-1})$ suffices which shows that all real-valued poles (and zeroes) appear in pairs $\{d_i,\ d_i^{-1}\}$ (c.f. Figure 1).
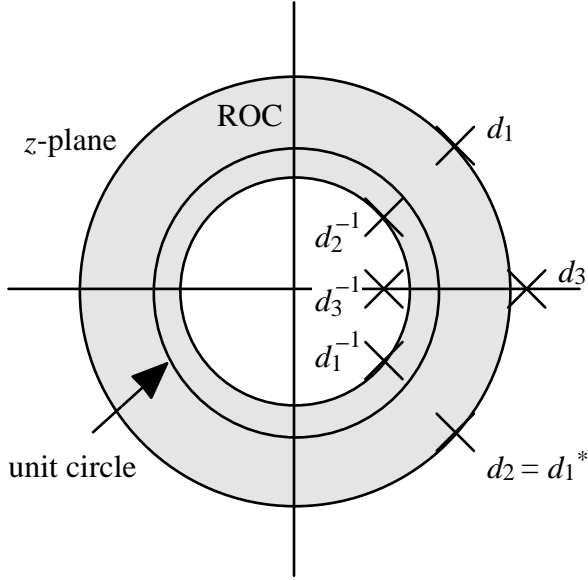

Figure 1: Graphical representation of the *z*-transform of a stable ($|z|=1$ in region-of-convergence) recursive filter with a real and even transfer function.

3.1 Series interconnection

A real and even transfer function $H(z)$ can be decomposed into two stable subsystems: a causal one $H_+(z)$ and an anti-causal one $H_-(z)$. The interconnection of the subsystems can be either in series $H_s(z) = H_{s+}(z) \cdot H_{s-}(z)$ or in parallel $H_p(z) = H_{p+}(z) + H_{p-}(z)$. Note that the poles of H_+ lie inside the unit circle and that the poles of H_- lie outside the unit circle. To minimize the computational complexity to N multiplications per subsystem, we demand that each subsystem should be a purely recursive system, i.e. no polynomial in z in the numerator of the transfer function is allowed. For the series interconnection we find

$$H_s(z) = \prod_{i=1}^{N} \frac{d_i - 1}{d_i - z^{-1}} \cdot \prod_{i=1}^{N} \frac{1 - d_i}{z - d_i} \equiv H_{s+}(z) \cdot H_{s-}(z) \quad (3)$$

For the parallel interconnection scheme we find

$$H_p(z) = H_{p+} + H_{p-} \equiv \frac{1}{2}\left(\prod_{i=1}^{N} \frac{d_i - 1}{d_i - z^{-1}} + \prod_{i=1}^{N} \frac{1 - d_i}{z - d_i} \right) \quad (4)$$

The overall parallel transfer function $H_p(z)$ has $2N$ poles (i.e. the poles of the two subsystems) and N zeroes. These zeroes in the *z*-transform are a function of the poles $\{d_i\}$ and hamper the quality of the Gaussian approximation, i.e., less suppression of the high frequencies. Deriche employs a parallel interconnection scheme of two full N^{th} order subsystems ($2N$ multiplications per subsystem) consisting of N poles and N zeroes. The zeroes supply the necessary degrees of freedom to allow an accurate approximation to the true Gaussian in exchange for twice the number of multiplications. To minimize the computational complexity we choose the series interconnection of a stable causal and a stable anti-causal subsystem.

3.2 Discrete-time transfer function

In the remainder of this paper we focus on the series interconnection and drop the subscript *s*. The discrete time (space) systems are described by their *z*-transforms which can be decomposed into a product of two stable subsystems H_+ (causal system) and H_- (anti-causal system). Using $\{d_i\}$ for the poles in the *z*-plane we obtain

$$H(z) = H_+(z) \cdot H_-(z) \qquad (5)$$

with

$$H_+(z) = \prod_{i=1}^{N} \frac{d_i - 1}{d_i - z^{-1}} \qquad H_-(z) = (-1)^N \prod_{i=1}^{N} \frac{d_i - 1}{z - d_i}$$

The resulting discrete time transfer function $H(\Omega)$ for arbitrary N

$$H(\Omega) = H(z)\big|_{z = e^{j\Omega}} = \prod_{i=1}^{N} (-1)^N H_{d_i}(\Omega) \qquad (6)$$

with

$$H_{d_i}(\Omega) = \frac{(d_i - 1)^2}{(d_i - z^{-1})(z - d_i)}\bigg|_{z = e^{j\Omega}} = \frac{(d_i - 1)^2}{2 d_i \cos\Omega - 1 - d_i^2} \quad (7)$$

Each term $H_{d_i}(\Omega)$ consists of a causal as well as the corresponding anti-causal pole. Although the individual terms $H_{d_i}(\Omega)$ are complex-valued for a complex pole $\{d_i\}$, the overall discrete time (space) transfer function is always real-valued because the complex-valued poles $\{d_i\}$ appear in conjugate pairs (eg. (2)). The contribution of a conjugate pair of poles to the discrete time Fourier transform is real-valued and shown below.

$$H_{d_i}(\Omega) H_{d_i^*}(\Omega) = \frac{\left((\operatorname{Re} d_i - 1)^2 + \operatorname{Im}^2 d_i\right)^2}{4 \operatorname{Im}^2 d_i (\cos\Omega - \operatorname{Re} d_i)^2 +} \qquad (8)$$
$$\left(\operatorname{Im}^2 d_i - (\cos\Omega - \operatorname{Re} d_i)^2 - \sin^2 \Omega\right)^2$$

3.3 Second moment of the discrete-time recursive filter

The second moment of a true Gaussian filter is equal to its variance σ^2. Using the derivative property of the discrete-time Fourier transform and the evenness of the Gaussian, we can show that the second moment in the time (space) domain corresponds to the second derivative of the transfer function at $\Omega=0$.

$$\mathrm{var}(g[n;\sigma]) = \sigma^2$$

$$= \sum_{n=-\infty}^{+\infty} n^2 g[n;\sigma] \quad \xleftrightarrow{\ \mathcal{F}\ } \quad -\frac{\partial^2 G(\Omega;\sigma)}{\partial \Omega^2}\bigg|_{\Omega=0} \qquad (9)$$

The second derivative of our recursive discrete-time transfer function at $\Omega=0$ is a measure for the filter's second moment (scale). Using eqs (6) and (7) we can show that

$$\sum_{n=-\infty}^{+\infty} n^2 h[n] = -\frac{\partial^2 H(\Omega)}{\partial \Omega^2}\bigg|_{\Omega=0} = \sum_{i=1}^{N} \frac{2d_i}{(d_i-1)^2} \qquad (10)$$

The equation above always yields a real-valued variance because the complex-valued poles d_i always appear in conjugate pairs (eg. (2)).

3.4 Implementation: difference equations

The cascade of both systems as illustrated in Figure 2 with the proper choice of $\{d_i\}$ can approximate the discrete Gaussian filter.
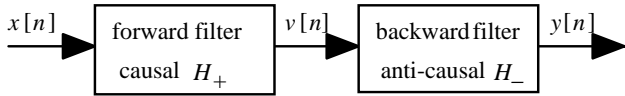

Figure 2: Two systems in series: one causal (H_+) and one anti-causal (H_-).

The subsystems H_+ and H_- correspond to a forward and a backward difference equation with real-valued coefficients $\{b_i\}$.

$$v[n] = \alpha x[n] - \sum_{i=1}^{N} b_i\, v[n-i] \quad \xleftrightarrow{z} \quad H_+(z) = \frac{\alpha}{1+\sum\limits_{i=1}^{N} b_i z^{-i}}$$

$$y[n] = \alpha v[n] - \sum_{i=1}^{N} b_i\, y[n+i] \quad \xleftrightarrow{z} \quad H_-(z) = \frac{\alpha}{1+\sum\limits_{i=1}^{N} b_i z^{i}} \qquad (11)$$

To guarantee that $H(\Omega=0) = 1$ which is equivalent to $H(z=1) = 1$ we write

$$\alpha = 1+\sum_{i=1}^{N} b_i z^{-i}\bigg|_{z=1} = 1+\sum_{i=1}^{N} b_i \qquad (12)$$

Note that the filter coefficients of the forward and backward filters are identical. The coefficients $\{b_i\}$ can be expressed as a function of the complex pole coordinates in the z-plane $\{d_i\}$. Since the two sub-filters are identical, we give the coefficients $\{b_i\}$ as a function of the anti-causal poles, i.e. the poles outside the unit circle. For arbitrary N the coefficients $\{b_i\}$ are:

$$b_N = (-1)^N b$$

$$b_{N-1} = (-1)^{N-1} b \sum_{i=1}^{N} d_i$$

$$b_{N-2} = (-1)^{N-2} b \sum_{i=2}^{N} \sum_{j=1}^{i-1} d_i\, d_j$$

$$b_{N-3} = (-1)^{N-3} b \sum_{i=3}^{N} \sum_{j=2}^{i-1} \sum_{k=1}^{j-1} d_i\, d_j\, d_k \qquad (13)$$

$$\vdots$$

$$b_{N-m} = (-1)^{N-m} b \sum_{i_m=m}^{N} \cdots \sum_{i_2=2}^{i_3-1} \sum_{i_1=1}^{i_2-1} d_{i_m} \cdots d_{i_2} d_{i_1}$$

$$\vdots$$

$$b_0 = (-1)^0 b\, d_1\, d_2 \cdots d_N \equiv 1$$

with

$$b \equiv (d_1\, d_2 \cdots d_N)^{-1} = \prod_{i=1}^{N} d_i^{-1}$$

Given the above closed-form relations between the filter coefficients $\{b_i\}$ and the z-poles $\{d_i\}$ we now have to find the set of poles that minimizes the error between the discrete recursive filter and the desired Gaussian.

4. Approximation of the discrete-time transfer function

The discrete-time transfer function $H(\Omega)$ for arbitrary order N can be constructed using eqs. (6) to (8). For $N = 3$, 4, and 5 we have calculated the poles $\{d_i\}$ that minimize the error between $H(\Omega)$ and the true Gaussian transfer function $G(\Omega; \sigma=2)$. We choose $\sigma = 2$ because it yields a Gaussian that can be considered bandlimited. We minimized two error criteria, the root square error (L^2 norm) and the maximum error (L^∞ norm), using a conjugate gradient iterative minimization procedure.

$$L^2 = \sqrt{\frac{1}{2\pi}\int_{-\pi}^{+\pi} (G(\Omega;\sigma) - H(\Omega))^2\, d\Omega}$$

$$= \sqrt{\sum_{n=-\infty}^{+\infty} (g[t;\sigma] - h[n])^2}$$

$$L^\infty = \lim_{k\to\infty} \sqrt[k]{\int_{-\pi}^{+\pi} (G(\Omega;\sigma) - H(\Omega))^k\, d\Omega} \qquad (14)$$

$$= \max_{\{|\Omega|<\pi\}} (\|G(\Omega;\sigma) - H(\Omega)\|)$$

The z-poles $\{d_i\}$ and residual errors for the Gaussian filters are listed in Table 1. The residual errors for both error criteria (L^∞ and L^2) are roughly the same. We notice a decrease of all errors as the filter order increases.

Table 1: The poles $\{d_i\}$ and residual errors (L^2, L^∞) for recursive Gaussian filters ($\sigma = q = 2$; var($h[n;q=2] = 4$) of order N=3, 4, and 5. The poles d_{2i} are given by $d_{2i} = d_{2i-1}^*$.

	N=3 L^2	N=3 L^∞	N=4 L^2	N=4 L^∞	N=5 L^2	N=5 L^∞
Re d_1	1.41650	1.40098	1.13228	1.12075	0.86430	0.85480
Im d_1	1.00829	1.00236	1.28114	1.27788	1.45389	1.43749
Re d_3	1.86543	1.85132	1.78534	1.76952	1.61433	1.61231
Im d_3	-	-	0.46763	0.46611	0.83134	0.82053
Re d_5	-	-	-	-	1.87504	1.87415
L^∞	12.7 E-3	12.0 E-3	4.57 E-3	4.32 E-3	1.79 E-3	1.64 E-3
L^2	7.17 E-3	7.21 E-3	2.42 E-3	2.48 E-3	0.91 E-6	0.95 E-6

The improvements in accuracy become clear when comparing these residual errors for N=3 with the residual errors of our previous third-order recursive filter [9] of scale q=1 (which has the same variance as a true Gaussian filter of scale σ=2.068): $L^\infty = 82.0$ E-3, $L^2 = 59.6$ E-3.

4.1 Gaussian derivative filters

Sample *differences* are used as approximations to the true derivatives.

$$\tfrac{1}{2}\left(h[n+1]-h[n-1]\right) \quad\xleftrightarrow{\mathcal{F}}\quad j\sin(\Omega)\,H(\Omega)$$

$$h[n-1]-h[n]+h[n+1] \xleftrightarrow{\mathcal{F}} -2(1-\cos(\Omega))\,H(\Omega) \tag{15}$$

These approximations are very accurate for the low frequencies. As a consequence, the derivative filters gain accuracy with increasing scale, σ. The implementation is given in [9]. We designed dedicated Gaussian approximations for the first and second Gaussian derivative filters by minimizing a Ω-weighted error between the new filters and $G(\Omega;\sigma=2)$. We use a Ω-weighted error instead of the weights corresponding to the sample differences from eq. (15) to allow subsequent scaling of the underlying Gaussian approximation. The dedicated filters are denoted by $H_{d,N}(\Omega)$ in which N denotes the filter order and d the derivative order for which the filter was designed. For a first derivative the error is weighted by Ω over the range $-\pi < \Omega < \pi$. For a second derivative the weight is Ω^2 over the same range. Table 2 lists the Ω^d-weighted residual errors for the first and second Gaussian derivative filters. Here we minimized the

Ω^d-weighted maximum error L_d^∞ (but computed the resulting Ω^d-weighted root square error L_d^2 as well)

$$L_d^\infty = \max_{\{|\Omega|<\pi\}}\left\{\left|\Omega^d\left(G(\Omega;\sigma)-H_{d,N}(\Omega)\right)\right|\right\}$$

$$L_d^2 = \sqrt{\tfrac{1}{2\pi}\int_{-\pi}^{+\pi}\Omega^d\left(G(\Omega;\sigma)-H(\Omega)\right)^2 d\Omega} \tag{16}$$

The improvements in accuracy become clear when comparing the residual errors with the residual errors of our previous recursive filter [9] which are given in Table 3.

Table 2: The poles $\{d_i\}$ and residual errors for recursive first and second Gaussian-derivative filters ($\sigma = 2$; var($h[n]$ $= 4$) of order N=3, 4, and 5. Note that $d_{2i} = d_{2i-1}^*$.

	first Gaussian derivative filters		
	$H_{1,3}$	$H_{1,4}$	$H_{1,5}$
Re d_1	1.31553	1.04185	0.77934
Im d_1	0.97057	1.24034	1.41423
Re d_3	1.77635	1.69747	1.50941
Im d_3	-	0.44790	0.80828
Re d_5	-	-	1.77181
L_d^∞	14.37 E-3	5.42 E-3	2.20 E-3
L_d^2	9.39 E-3	3.31 E-3	1.34 E-3

	second Gaussian derivative filters		
	$H_{2,3}$	$H_{2,4}$	$H_{2,5}$
Re d_1	1.22886	0.94570	0.69843
Im d_1	0.93058	1.21064	1.37655
Re d_3	1.70493	1.60161	1.42631
Im d_3	-	0.42647	0.77399
Re d_5	-	-	1.69668
L_d^∞	18.81 E-3	7.25 E-3	3.11 E-3
L_d^2	13.96 E-3	4.89 E-3	2.01 E-3

Table 3: For comparison: the Ω^d-weighted residual errors of the recursive Gauss, first and second derivative filters presented by [9].

σ=2.068 q=1	d=0; N=3 H	d=1; N=3 H	d=2; N=3 H
L_d^∞	82.10 E-3	90.80 E-3	118.62 E-3
L_d^2	42.23 E-3	52.44 E-3	83.88 E-3

Comparing the residual errors for N=3 (Table 2) with the residual errors of our previous third-order recursive filter (Table 3) (which has the same variance as a true Gaussian filter of scale $\sigma = 2.068$) we notice an improvement in accuracy of a factor of 5 in maximum error. For N=5 the maximum error becomes roughly 40 times smaller.

4.2 Scaling the discrete-time filter

In order to create a recursive Gaussian filter of arbitrary scale, we need to find an analytical relation between the z-poles $\{d_i\}$ and a specified scaling parameter q. In our recursive filter q plays the role of σ in the true Gaussian filter of eq. (1). Scaling of a continuous-time filter corresponds to inverse scaling of its Laplace transform

$$h\left(\frac{t}{q}\right) \quad \overset{\mathcal{L}}{\longleftrightarrow} \quad |q|H(s\,q) \tag{17}$$

Scaling of a continuous transfer function is achieved by scaling the s-poles $\{s_i\}$ parallel as well as perpendicular to the $j\omega$-axis ($s_i \rightarrow s_i\,q^{-1}$). Such a scaling property does not exist in the discrete domain, but can be borrowed from the continuous domain. A discrete filter is equivalent to a continuous-time filter that consists solely of weighted impulses as given by $\Sigma\ h(nT)\bullet\delta(t–nT)$. The Laplace transform of this "sampled" continuous filter is periodic in ω with a period $2\pi/T$. The geometric transform that relates the discrete filter's z-transform to the sampled continuous filter's Laplace transform ([14] chapter 10.8.1) is

$$z = e^{s\,T} \tag{18}$$

This geometric transform maps one period of the s-plane into the z-plane. It can also be used to translate the scaling property to the discrete domain. By analogy to scaling in the continuous domain, scaling of the z-poles should occur both in parallel (phase of z-poles) and perpendicular (modulus of z-poles) to the unit circle in the z-plane (The unit circle is the locus of frequency in the Fourier transform).

$$d_i \rightarrow e^{\frac{\log(|d_i|)}{q}}\,e^{j\frac{\arctan(\operatorname{Im}d_i/\operatorname{Re}d_i)}{q}} = |d_i|^{\frac{1}{q}}\,e^{j\frac{\theta(d_i)}{q}} \tag{19}$$

with $\theta(d_i) = \arctan(\operatorname{Im}d_i/\operatorname{Re}d_i)$. Note that we scale relative to the approximation to a Gaussian of $\sigma_0 = 2$. The exact variance of the scaled filter still deviates from the specified scale of $q\,\sigma_0$. The resulting variance of the scaled filter can be obtained by combining eqs (10) and (19)

$$\operatorname{var}(h[n]) = \sum_{i=1}^{N} 2\,|d_i|^{\frac{1}{q}}\,e^{j\frac{\theta(d_i)}{q}}\left(|d_i|^{\frac{1}{q}} - e^{j\frac{\theta(d_i)}{q}}\right)^{-2} \tag{20}$$

A few iterations of a simple linear extrapolation scheme finds the value for q that yields a filter $h[n]$ with the specified variance.

5. Results

In previous work [9] we used the continuous-time Abramowitz approximation to the Gaussian. The backward/forward difference method is used to transform this continuous-time filter into a discrete-time filter. The backward/forward difference method only works for very-low frequencies, i.e., for large scales $\sigma > 10$. The method yields poor Gaussian approximations for small scales. Now we can compare the accuracy of our previous method [9] with our new discrete-time design method. As depicted in the top row of Figure 4 we notice that the resulting filter becomes more accurate as q increases. There remains a difference because the continuous-time Abramowitz approximation H_A

$$H_A(\omega) = \frac{1}{a_0 + a_2\omega^2 + a_4\omega^4 + a_6\omega^6}$$

with

$$\operatorname{var}(h(t)) = -\frac{\partial^2 H_A(\omega)}{\partial\omega^2}\bigg|_{\omega=0} = \frac{2a_2}{a_0}\bigg|_{\substack{a_0=2.4909\\a_2=1.4660}} = 1.177 \tag{21}$$

has a 1.177 times larger variance than the corresponding true Gaussian. The new recursive filter works well over a wide range of scales, from σ=1 to σ =100.
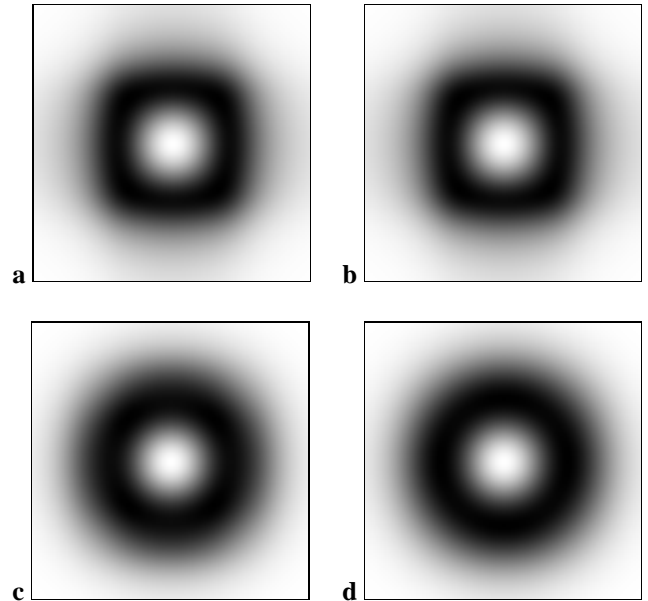

Figure 3: Laplacian transfer functions with $\sigma = 2$: **a)** previous recursive Laplacian-of-Gaussian [9]; **b)** new third-order recursive filter $H_{2,3}$; **c)** new fifth-order recursive filter $H_{2,5}$; **d)** true Laplacian-of-Gaussian transfer function.
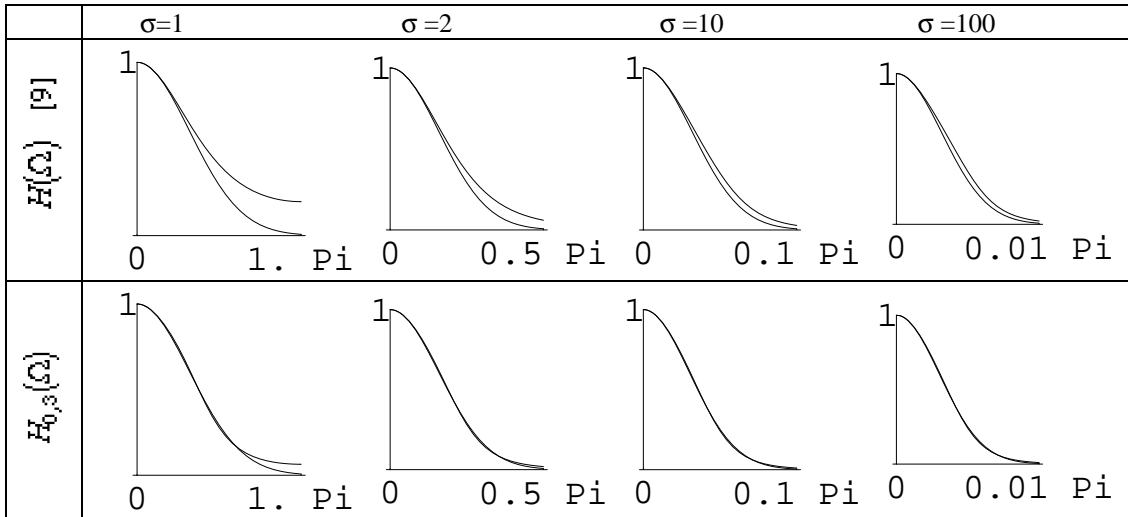
Figure 4: Discrete-time transfer functions: true Gaussian transfer function $G(\Omega;\sigma)$ and $H(\Omega)$ of *equal* variance. The true Gaussian is the lower of the two. **top)** the previous recursive filter [9]; **bottom)** the new recursive filter $H_{0,3}(\Omega)$.

5.1 Isotropy

The need for a better approximation to the Gaussian derivative transfer function comes from isotropy requirements. First and second order Gaussian derivatives yield some well-known rotation invariant operators:

$$|\nabla g| = \sqrt{g_x^2 + g_y^2} \qquad \nabla^2 g = g_{xx} + g_{yy} \qquad (22)$$

Our previous recursive Gaussian derivative filters produced a non-isotropic response. Figure 3 shows that the new fifth-order recursive filter yields an excellent approximation to an isotropic response.

6. Conclusions

Our recursive Gaussian derivative filters of order N (N=3,4,5) require only $2N$ multiplications per pixel (per dimension). They offer a constant execution time independent of the Gaussian scale parameter σ. The accuracy of our approximation is high and increases with increasing filter order. To obtain an excellent isotropic response in 2-D images we advise a third-order Gaussian filter ($H_{0,3}$), a fourth-order first derivative filter ($H_{1,4}$), and a fifth-order second derivative filter ($H_{2,5}$). The exact inverse filters are trivial to obtain. In linear scale space the diffusion equation defines the image at the next scale. However, the results at all scales can be computed directly from the original image by Gaussian blurring in constant time.

7. Acknowledgments

8. References

[1] A. Witkin, "Scale-space filtering," Proc. Int. Joint Conf. on Artificial Intel., Karlsruhe, Germany, 1019-1021, 1983.

[2] J.J. Koenderink, "The structure of images," *Biological Cybernetics*, vol. 50, 363-370, 1984.

[3] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, 679–698, 1986.

[4] P.W. Verbeek and L.J. van Vliet, "On the localization error of curved edges in low-pass filtered 2-D and 3-D images," *IEEE-PAMI*, vol. 16, 726-733, 1994.

[5] L.J. van Vliet, *Grey-Scale Measurements of Multi-Dimensional Digitized Images*, Delft Univ. Press, 1993.

[6] L. Haglund, "Adaptive multidimensional filtering," Linköping University, Sweden, 1992.

[7] L.J. van Vliet and P.W. Verbeek, "Estimators for orientation and anisotropy in digitized images," Proc. ASCI'95, Heijen (Netherlands), 442-450, 1995.

[8] A. Papoulis, *Signal Analysis*. McGraw–Hill, 1977.

[9] I.T. Young and L.J. van Vliet, "Recursive implementation of the Gaussian filter," *Signal Proc.*, vol 44, 139-151, 1995.

[10] M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*, 1965.

[11] R. Deriche, "Separable recursive filtering for efficient multi–scale edge detection," Proc. Int. Workshop Machine Vision Machine Intelligence, Tokyo, Japan, 18–23, 1987.

[12] R. Deriche, "Fast algorithms for low–level vision," *IEEE Trans. Pat. Anal. Mach. Intel.*, vol. 12, 78–87, 1990.

[13] R. Deriche, "Recursively implementing the Gaussian and its derivatives," Proc. 2nd International Conference on Image Processing, Singapore, 263-267, 1992.

[14] Oppenheim, A.S. Willsky, and I.T. Young, *Signals and Systems*. Englewood Cliffs, NJ: Prentice–Hall, 1983.