Spatio-Frequency Analysis of Scale-Space Filtering

Luc Florack

Abstract

We study implementation issues for spatial convolution filters and their Fourier alternative, with the aim to optimise the accuracy of filter output. We focus on Gaussian scale-space filters and show that there exists a trade-off scale that subdivides the available scale range in the sense that the smaller and larger scales are best extracted by Fourier, respectively spatial filtering.

1 Introduction

A Gaussian scale-space representation of a raw image f is obtained by convolution with a normalised, appropriately scaled Gaussian filter ϕ:

$$u(\mathbf{x};\sigma) = \int f(\mathbf{y})\,\phi(\mathbf{x}-\mathbf{y};\sigma)\,d\mathbf{y}, \tag{1}$$

with, in n dimensions,

$$\phi(\mathbf{x};\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}^{\,n}} \exp\left(-\frac{1}{2}\frac{\|\mathbf{x}\|}{\sigma^2}\right) . \tag{2}$$

Higher order derivatives require similar filterings based on derivatives of the zeroth order Gaussian. Note that a scale-space representation satisfies the diffusion equation

$$u_t = \Delta u\,, \tag{3}$$

in which the evolution parameter is related to scale by $\sigma = \sqrt{2t}$.

There are various options for implementating Eq. (1) in practice. Lindeberg considers discrete versions of the diffusion equation, maintaining a continuous scale parameter, which has the advantage that certain scale-space axioms are manifestly preserved in the discrete setting [1, 2]. Other approaches are based on a discrete implementation of the solution instead. Among these, three techniques are particularly noteworthy. Deriche describes a recursive scheme, which in the case of scale-space filtering has the advantage that computation time is scale-independent [3, 4]. Another way to achieve scale-independent computation time is to consider Eq. (1) in frequency space, based on the well-known fact that convolution in the spatial domain becomes pixelwise multiplication in the Fourier domain. Of course it requires a two-fold mapping of the raw image to and from Fourier space. Note that we do *not* need to map filters if we know their Fourier forms analytically, and that if we are interested in several levels of scale and/or multiple partial derivatives, the forward transform of the raw image needs to be carried out only once. The third and most straightforward method is a direct discretisation of Eq. (1) as it stands, in which case computation times will scale with effective filter size.

In this article we investigate the latter two options, and pose the following questions:

- The criterion being *accuracy* of filter output, should one use spatial convolution or multiplicative Fourier filters?
- If the answer depends on scale, what is the quantitative criterion for deciding between these two domains?

	condition	"synthesis" or inverse FT	"analysis" or forward FT
CFT: $t, \Omega \in \mathbb{R}$	$\int dt \, \lvert g(t) \rvert < \infty$	$g(t) \stackrel{\text{def}}{=} \frac{1}{2\pi} \int d\Omega \, G(\Omega) \, e^{i\Omega t}$	$G(\Omega) = \int dt \, g(t) \, e^{-i\Omega t}$
DFT: $n, k = 0, \ldots, N-1$	☺	$g_n^N \stackrel{\text{def}}{=} \frac{1}{N} \sum_{k=0}^{N-1} G_k^N \, e^{2\pi i n k / N}$	$G_k^N = \sum_{n=0}^{N-1} g_n^N \, e^{-2\pi i n k / N}$
DSFT$_1$: $n \in \mathbb{Z}, \omega \in [-\pi, \pi]$	$\sum_{n \in \mathbb{Z}} \lvert g_n^\infty \rvert < \infty$	$g_n^\infty \stackrel{\text{def}}{=} \frac{1}{2\pi} \int_{-\pi}^{\pi} d\omega \, G^\infty(\omega) \, e^{i\omega n}$	$G^\infty(\omega) = \sum_{n \in \mathbb{Z}} g_n^\infty \, e^{-i\omega n}$
DSFT$_2$: $k \in \mathbb{Z}, \tau \in [-\pi, \pi]$	$g^\infty \in C^0\left([-\pi, \pi]\right)$	$g^\infty(\tau) \stackrel{\text{def}}{=} \frac{1}{2\pi} \sum_{k \in \mathbb{Z}} G_k^\infty \, e^{i\tau k}$	$G_k^\infty = \int_{-\pi}^{\pi} d\tau \, g^\infty(\tau) \, e^{-i\tau k}$

Table 1: The various 1D Fourier transform pairs used in this article. The only one operationally defined is DFT/FFT.

2 Theory

2.1 Various Fourier Transforms

In continuous formulations of scale-space theory one frequently considers an analytical Fourier transform in which image representations are modelled as *analogue* functions in both spatial as well as frequency domains. This type of Fourier transform, which exists only on paper, will henceforth be referred to by the acronym CFT (" Continuous Fourier Transform"). In practice one deals with the Discrete Fourier Transform (DFT)—or, if conditions permit, the optimised Fast Fourier Transform (FFT), *cf.* Cochran *et al.* [5]. It is therefore of interest to study the precise relation between CFT and DFT/FFT. In order to do that it is useful to distinguish a third type of Fourier Transform, applicable to infinite discrete sequences, *viz.* the Discrete Sequence Fourier Transform (DSFT), of which there are essentially two types, depending on which of the domains one chooses to be the discrete one. The formulas for the distinct Fourier transform pairs are listed in Table 1 for the 1D case, including conditions for sufficiency. Further details are given below.

- CFT is appropriate for an *analogue* function $g(t)$ defined on all of $t \in \mathbb{R}$. If $g \in L^1(\mathbb{R})$ then $G(\Omega)$ is continuous and bounded.
- DFT relates to a *finite* sequence of discrete samples[1] g_n^N, say $n = 0, \ldots, N-1$. DFT is *operationally well-defined*, *cf.* Numerical Recipes [6].
- DSFT$_1$ applies to discrete sequences of *infinite* length, and is sort of "in-between", which is reflected by the lack of symmetry of the inverse expression. If the sequence g_n^∞ is absolutely convergent, *i.e.* if $\sum_{n \in \mathbb{Z}} \lvert g_n^\infty \rvert < \infty$, then the Fourier series will be uniformly convergent to a continuous, 2π-periodic function $G^\infty(\omega)$. Note that this precludes nontrivial periodic sequences g_n^∞. The Fourier spectrum is of an analogue nature, and is band-limited.
- The same reasoning holds for DSFT$_2$, but with the role of function domains interchanged.

All listed Fourier pairs are one-to-one. Thus we cannot expect to gain accuracy merely by switching domains using any of the given transforms. For example, a discrete model g_n^N of a continuous filter $g(t)$ which is very narrow, a few pixels say, will suffer from severe discretisation artifacts in the temporal domain. DFT will yield a Fourier filter G_k^N hampered by severe aliasing effects.

This argument suggests that, as far as accuracy is concerned, it makes no difference which domain one chooses to work in. This conclusion is, however, premature. The trick we can exploit in order to optimise accuracy of computation (disregarding speed requirements) lies in the *vertical* relations of Table 1: The discretisation of an analytically given CFT differs fundamentally from the DFT of a discretised spatial filter. We study this in detail below.

[1] The superscript N is not an exponent, but serves as a reminder.

parameter glossary		
relative to grid	$\rho \stackrel{\text{def}}{=} \frac{T}{\sigma}$	$\eta \stackrel{\text{def}}{=} \sigma F$
relative to scope	$\delta \stackrel{\text{def}}{=} \frac{t_c}{\sigma}$	$\zeta \stackrel{\text{def}}{=} \sigma \Omega_c$
grid-scope relation	$t_c \stackrel{\text{def}}{=} \frac{1}{2} NT$	$\Omega_c \stackrel{\text{def}}{=} \frac{1}{2} MF$
grid-grid relation	$N \stackrel{\text{def}}{=} M$	
physical constraint (uncertainty principle)	$t_c = \frac{\pi}{F}$	$\Omega_c \stackrel{\text{def}}{=} \frac{\pi}{T}$
other dependencies	$\delta = \frac{1}{2} N\rho$	$\zeta = \frac{1}{2} M\eta$

Table 2: A glossary of scale-related parameters and mutual dependencies. The ones marked with "def" can be regarded as the independent ones.

2.2 Optimising Accuracy of Filter Output

Again, without loss of generality, we consider the 1D case. We refer to Table 1 for the notation. We start with an analytical model of a filter $g(t)$, *e.g.* a normalised Gaussian or one of its derivatives, and assume its Fourier representation $G(\Omega)$ is known in analytically closed form. Let us say we sample the temporal filter, *i.e.* we take

$$g_n^\infty \stackrel{\text{def}}{=} g(nT) \quad n \in \mathbb{Z}\,, \tag{4}$$

in which we have introduced a global constant T, the *temporal sampling interval*, or "pitch" for brevity. Clearly this is only one step towards an implementation; we will also have to decide on a finite *number of samples*, thereby introducing a second system parameter N (recall Footnote 1):

$$g_n^N \stackrel{\text{def}}{=} g(nT) \quad n \in \{0, \ldots, N-1\}\,. \tag{5}$$

With this choice of index range, we assume that $g(t) = \phi(t - t_c)$, shifting the base point to the centre: $t_c = \frac{1}{2}NT$. (The global constants are incorporated in the notation only if it might cause confusion otherwise.)

We want to know what it means to say that Eq. (5) is a discrete approximation of $g(t)$, and to which extent it is. For simplicity we disregard errors induced by filter truncation. This does not affect the main result obtained below, since truncation effects pertain to scale limitations at the boundaries of the physical scale range, whereas subsequent considerations will turn out to affect intermediate scales only.

Table 2 contains a glossary of parameters that will be introduced in the next few sections. Please refer to this table for explanations and interrelations of symbols introduced below.

2.3 Discretisation in the Temporal Domain

We disregard truncation to isolate the effect induced by discrete sampling, Eq. (4), and consider temporal filtering.

The sequence g_n^∞ has an equivalent Fourier transform $G^\infty(\omega)$ (DSFT_1), which we can think of as a 2π-periodic cycle. The following equation relates DSFT_1 to CFT (recall Eq. (4) and Table 1):

$$G^\infty(\Omega T) = \frac{1}{T} \sum_{r \in \mathbb{Z}} G\left(\Omega + \frac{2\pi r}{T}\right)\,. \tag{6}$$

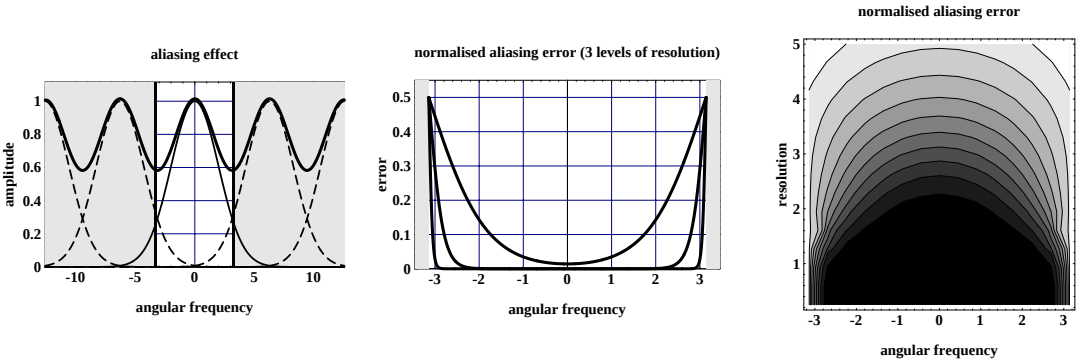

Figure 1: The left graph illustrates the aliasing effect. The thin curve corresponds to the ideal profile, while the thick curve follows the actual one hampered by superposition of spurious copies indicated by the dashed lines. The middle figure shows the quantitative normalised aliasing error (explained in the text) as a function of angular frequency $\omega \in [-\pi, \pi]$ for the normalised Gaussian filter at three levels of temporal resolution, $\rho = 0.5, 1.0, 2.0$ respectively. The contour plot on the right shows the same error in the (ω, ρ)-plane. Details are given by Eqs. (10–11).

Note the differences in dimensionality of the various quantities involved.

To prove Eq. (6), consider Table 1, from which it follows that

$$g_n^\infty \stackrel{\text{def}}{=} g(nT) = \frac{1}{2\pi} \int d\Omega\, G(\Omega)\, \exp(in\Omega T)\,.$$

Substituting $\Omega T \equiv \omega'$ and splitting up the integral into a sum of integrals over 2π-intervals yields

$$g_n^\infty = \frac{1}{2\pi T} \sum_{r\in\mathbb{Z}} \int_{(2r-1)\,\pi}^{(2r+1)\,\pi} d\omega'\, G\left(\frac{\omega'}{T}\right)\, \exp(i\omega' n)\,.$$

Substituting variables once more, setting $\omega' = \omega + 2\pi r$, and using the fact that $\exp(2\pi i n r) = 1$ for all $r \in \mathbb{Z}$, we obtain

$$g_n^\infty = \frac{1}{2\pi T} \sum_{r\in\mathbb{Z}} \int_{-\pi}^{\pi} d\omega\, G\left(\frac{\omega + 2\pi r}{T}\right)\, \exp(i\omega n)\,.$$

Interchanging sum and integral in this result, and comparing it with the DSFT_1-inverse entry,

$$g_n^\infty = \frac{1}{2\pi} \int_{-\pi}^{\pi} d\omega\, G^\infty(\omega)\, \exp(i\omega n)\,,$$

establishes the proof. (The result is evaluated at $\omega = \Omega T$.)

Eq. (6) relates the Fourier transform of the sampled filter (l.h.s.) to that of its ideal model (r.h.s.), which is illustrated in Fig. 1. We observe the following.

- We must relate physical frequency $\Omega \in \mathbb{R}$ to dimensionless frequency $\omega \in [-\pi, \pi]$ by $\omega \equiv \Omega T$. The physical interval comprises angular frequencies $\Omega \in [-\pi/T, \pi/T]$. The half-width of the physical frequency interval is well-known as the *Nyquist critical frequency*. Its numeric value depends on parametric convention; typically one encounters $\Omega_c = \pi/T$, $F_c = 1/(2T)$, $\omega_c = \pi$, or $f_c = 1/2$, in the various self-explanatory frequency parametrisations.

- On the physical interval the ideal filter $g(t)$ is never perfectly realized by the sequence $g_n^\infty = g(nT)$ due to the well-known phenomenon of *frequency aliasing*, the accumulation of erroneous periodic copies of $G(\Omega)$ when analysed in Fourier space (all $r \neq 0$ terms on the r.h.s.).

- The absolute error depends on the decay of $G(\Omega)$, more precisely, on its values outside the physical interval.

Moreover we conclude that we can always make the aliasing error negligible—at least if we keep away from the physical boundaries $\Omega = \pm\Omega_c$—by taking T suitably small relative to filter width. In practice the input data limit the possibility of grid refinement, and we shall assume henceforth that input data and filters are represented on identical grids. (We could then set $T \equiv 1$, measuring everything in terms of grid units, but for the sake of argument below we won't do that.)

The aliasing error depends on frequency Ω and on the sampling interval T. Of course it depends on any additional parameter attached to the filter, notably scale, differential order, *etc.* We can quantify the error as follows. Let us define the aliasing error as the ratio of spurious terms ($r \neq 0$) and main term ($r = 0$) as defined on the r.h.s. of Eq. (6):

$$\varepsilon_{\text{FA}}(\Omega; T) \stackrel{\text{def}}{=} \left| \sum_{r \in \mathbb{Z} \setminus \{0\}} \frac{G\left(\Omega + \frac{2\pi r}{T}\right)}{G\left(\Omega\right)} \right| \tag{7}$$

for all $\Omega \in [-\Omega_c, \Omega_c]$. A normalised aliasing error can be defined as the ratio of spurious terms and actual terms ($r \in \mathbb{Z}$):

$$\bar{\varepsilon}_{\text{FA}}(\Omega; T) \stackrel{\text{def}}{=} \frac{\varepsilon_{\text{FA}}(\Omega; T)}{\varepsilon_{\text{FA}}(\Omega; T) + 1} \,. \tag{8}$$

The normalised aliasing error is confined to the unit interval: For all Ω and T we have $0 \leq \bar{\varepsilon}_{\text{FA}}(\Omega; T) \leq 1$. Note that both errors have only been defined on the physical interval, and that they always have appreciable magnitudes at the critical frequencies. *E.g.* for an even filter we have $\varepsilon_{\text{FA}}(\pm\Omega_c; T) \geq 100\%$, or $\bar{\varepsilon}_{\text{FA}}(\pm\Omega_c; T) \geq 50\%$.

For a normalised Gaussian filter with inner scale σ, *i.e.* Eq. (2) for the 1D case,

$$\phi(t; \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\frac{t^2}{\sigma^2}\right) , \tag{9}$$

the aliasing errors of Eqs. (7–8) are given by

$$\varepsilon_{\text{FA}}(\omega; \rho) = \varepsilon(\omega; \rho) \,, \tag{10}$$

$$\bar{\varepsilon}_{\text{FA}}(\omega; \rho) = \bar{\varepsilon}(\omega; \rho) \,, \tag{11}$$

with $\omega = \Omega T$ as usual, and with dimensionless resolution parameter ρ defined as the ratio of grid unit and inner scale: Recall Table 2. The functions $\varepsilon : [-\pi, \pi] \times \mathbb{R}^+ \to \mathbb{R}$ and $\bar{\varepsilon} : [-\pi, \pi] \times \mathbb{R}^+ \to [0, 1]$ are defined as follows:

$$\varepsilon(\theta; \lambda) \stackrel{\text{def}}{=} 2\sum_{k=1}^{\infty} \cosh\left(\frac{2\pi k\theta}{\lambda^2}\right) \exp\left(-\frac{2\pi^2 k^2}{\lambda^2}\right) , \tag{12}$$

$$\bar{\varepsilon}(\theta; \lambda) \stackrel{\text{def}}{=} \frac{\varepsilon(\theta; \lambda)}{\varepsilon(\theta; \lambda) + 1} \,. \tag{13}$$

Recall Fig. 1 for an illustration.

2.4 Discretisation in the Frequency Domain

In the previous section we have quantified approximation errors caused by sampling a filter in the *temporal* domain. Of course we can start out from the *frequency* representation. This leads to results very similar to the ones already obtained. We list them here without duplication of similar proofs.

Sampling a Fourier filter yields a sequence

$$G_k^{\infty} \stackrel{\text{def}}{=} G(kF) \quad k \in \mathbb{Z} \,, \tag{14}$$

in which F is the *frequency sampling interval*. Truncating its tails gives (recall Footnote 1)

$$G_k^{M} \stackrel{\text{def}}{=} G(kF) \quad k \in \{0, \ldots, M-1\} \,, \tag{15}$$

again assuming the centre of gravity of $G(\Omega)$ half-way the field of view.

Again ignoring filter truncation, the effect of discrete frequency sampling now requires us to study the relation between CFT and DSFT_2. Again recall Table 1, and Eq. (14):

$$g^{\infty}(tF) = \frac{1}{F} \sum_{s \in \mathbb{Z}} g\left(t + \frac{2\pi s}{F}\right) . \tag{16}$$

Apparently:

- Physical time relates to dimensionless time by $\tau = tF$; the physical interval is $t \in [-\pi/F, \pi/F]$. The half-width t_c of this interval is the temporal counterpart of the Nyquist critical frequency.
- On the physical interval we have *temporal aliasing*.
- The error is caused by the spurious wraparound of temporal filter tails propagating back into the physical interval.

The temporal aliasing error for a sampled Fourier filter depends on time t and frequency sampling interval F. Let us define it as the ratio of spurious terms ($s \neq 0$) and main term ($s = 0$) as defined on the r.h.s. of Eq. (16):

$$\varepsilon_{\mathrm{TA}}(t; F) \stackrel{\mathrm{def}}{=} \left| \sum_{s \in \mathbb{Z} \setminus \{0\}} \frac{g\left(t + \frac{2\pi s}{F}\right)}{g(t)} \right| \qquad t \in [-t_c, t_c] . \tag{17}$$

The normalised aliasing error can defined as the ratio of spurious terms and actual terms ($s \in \mathbb{Z}$):

$$\bar{\varepsilon}_{\mathrm{TA}}(t; F) \stackrel{\mathrm{def}}{=} \frac{\varepsilon_{\mathrm{TA}}(t; F)}{\varepsilon_{\mathrm{TA}}(t; F) + 1} \qquad t \in [-t_c, t_c] . \tag{18}$$

Again we must restrict ourselves to the physical time interval, and again we see that temporal aliasing always becomes significant near the interval boundaries: $\varepsilon_{\mathrm{TA}}(\pm t_c; F) \geq 100\%$, or $\bar{\varepsilon}_{\mathrm{TA}}(\pm t_c; F) \geq 50\%$ (for an even filter). For the case of Eq. (9) the temporal aliasing errors of Eqs. (17–18) are given by

$$\varepsilon_{\mathrm{TA}}(\tau; \eta) = \varepsilon(\tau; \eta) \quad \text{and} \tag{19}$$

$$\bar{\varepsilon}_{\mathrm{TA}}(\tau; \eta) = \bar{\varepsilon}(\tau; \eta) , \tag{20}$$

with $\tau = tF$, and with dimensionless scale parameter η defined as the product of grid unit and temporal inner scale: Table 2.

2.5 The Trade-Off Scale

The results sofar can be summarised as follows.

- Temporal sampling limits physical frequencies and causes frequency aliasing, while
- Frequency sampling limits the physical time scope and causes temporal aliasing.

The relevant error measures have been defined in Eqs. (10–11) and Eqs. (19–20), respectively, both of which are expressed in terms of the functions defined in Eqs. (12–13).

Let us now study the exact trade-off. We henceforth take identical grids in both domains (a DFT/FFT requisite: *Cf.* Table 1):

$$M \stackrel{\mathrm{def}}{=} N . \tag{21}$$

Given the dependencies of Table 2, it is convenient to express all results in terms of *two* free parameters: The number of samples N, and canonical scale

$$\lambda \stackrel{\mathrm{def}}{=} \log \frac{\sigma}{T} . \tag{22}$$

Note that a naive counting argument suggests a number of *three*: In Table 2 we encounter 11 parameters and 8 independent equations, *ergo* 3 unknowns, say N, σ, and T. However, for reasons of scale invariance, absolute values of "hidden scale parameters", for which we can take T in the case at hand, should always be discarded; only *dimensionless* variables count. Indeed, whether T represents a nanosecond or a billion years is completely immaterial. This argument naturally leaves us with the two degrees of freedom mentioned above.

The fundamental role of scale has been extensively discussed in the literature on scale-space theory (*cf.* existing books and the references therein [7, 8, 2, 9]), but what about N? It represents the number of degrees of freedom of our raw input data, and as such affects our image representation. We will see below that N naturally constrains the "depth" of a scale-space image in a trade-off with a prescribed accuracy. The parameter also shows up in the uncertainty principle in Table 2 after we discretise and truncate: $NTF = 2\pi$. This follows by fitting N temporal samples into the physical time range, $NT = 2t_c = 2\pi/F$, or, equivalently, by fitting the same number of frequency samples into the physical frequency range ($T \leftrightarrow F$, $t_c \leftrightarrow \Omega_c$). Dimensionless resolution ρ and dimensionless scale η are subject to the same uncertainty principle as the grid constants T and F, yet without reference to these underlying hidden scales: $N\eta\rho = 2\pi$. Note also that the product of full scopes in temporal and frequency domains equals 2π times the number of grid points N: $2t_c \times 2\Omega_c = 2\pi N$, the dimensionless counterpart of which is $2\delta \times 2\zeta = 2\pi N$, which gives us a physical interpretation of the number of degrees of freedom as the volume of the available space-frequency product space (divided by 2π). Thus the uncertainty principle can be stated in various ways, but the crux is invariably the same: It expresses a fundamental trade-off between *local detail in one domain*, and *global structure in the other*, *e.g.* temporal graininess versus frequency scope, temporal inner scale versus frequency outer scale, and so forth. The product of corresponding characteristic scales always exceeds a fundamental threshold.

Since the aliasing errors have been defined as relative errors in dynamic range, and apply to an overlapping range of scales, they are in principle comparable. Recalling Eqs. (7–8), and Eqs. (17–18), let us define the aliasing error balance as

$$\beta_N(\Omega, t; F, T) \stackrel{\text{def}}{=} \frac{\bar{\varepsilon}_{\text{FA}}(\Omega; T)}{\bar{\varepsilon}_{\text{TA}}(t; F)}\,, \tag{23}$$

with domain $(\Omega, t) \in [-\Omega_c, \Omega_c] \times [-t_c, t_c]$, and subject to the following relation between sampling constants and critical parameter values:

$$t_c = \frac{\pi}{F}\ ,\ \Omega_c = \frac{\pi}{T}\ ,\ NTF = 2\pi\,. \tag{24}$$

Combining Eqs. (10–11), Eqs. (19–20) and Eq. (23), we obtain the following aliasing error balance for the case of Eq. (9), defined on $(\omega, \tau) \in [-\pi, \pi] \times [-\pi, \pi]$:

$$\beta_N(\omega, \tau; \lambda) \stackrel{\text{def}}{=} \frac{\bar{\varepsilon}_{\text{FA}}(\omega; \rho)}{\bar{\varepsilon}_{\text{TA}}(\tau; \eta)} = \frac{\bar{\varepsilon}(\omega; \rho)}{\bar{\varepsilon}(\tau; \eta)}\,, \tag{25}$$

with $\omega = \Omega T$, $\tau = tF$, and with parameter dependencies given by Table 2 and Eq. (22). We can use this result to establish an accuracy criterion for deciding between temporal or Fourier implementation of our filters, at least if the effects of truncation can be ignored (the conditions of which need to be verified). The balance has been defined as a function of time and frequency, giving some leeway to semantical considerations. For example, one may exploit the (ω, τ)-dependency if one knows something about frequency characteristics and location of the objects of interest. In that case one may choose to optimise within a fiducial time-frequency window.

If we need to decide between the two options *in advance* we may not be able to anticipate the details of such tasks. It may also be the case that a certain task cannot easily be related to time-frequency characteristics, or that we have to face a plethora of different tasks using a standard implementation. In all these cases it may be useful to have an unbiased accuracy trade-off. To this end we may use suitable norms of the normalised aliasing errors to determine a rule of thumb. Recall Eq. (23) and Eqs. (10–11) and (25). For a 2π-periodic function $\bar{\varepsilon} : [-\pi, \pi] \to [0, 1]$ we define the norm as the expectation value presuming a uniform distribution on the unit circle:

$$\|\bar{\varepsilon}\| \stackrel{\text{def}}{=} \frac{1}{2\pi} \int_{-\pi}^{\pi} d\theta\, \bar{\varepsilon}(\theta)\,. \tag{26}$$

The global aliasing error balance for a zeroth order Gaussian can then be defined as the ratio of norms of $\bar{\varepsilon}_{\text{FA}}(\omega;\rho)$ and $\bar{\varepsilon}_{\text{TA}}(\tau;\eta)$:

$$\beta_N(\lambda) \stackrel{\text{def}}{=} \frac{\|\bar{\varepsilon}_{\text{FA}}(\,.\,;\rho)\|}{\|\bar{\varepsilon}_{\text{TA}}(\,.\,;\eta)\|} = \frac{\|\bar{\varepsilon}(\,.\,;\rho)\|}{\|\bar{\varepsilon}(\,.\,;\eta)\|}\,, \tag{27}$$

in which the norms are taken with respect to the first argument, and in which the parameters ρ and η are considered as functions of N and λ.

Using the balance $\beta_N(\lambda)$ of Eq. (27) as a criterion we may opt for

- temporal implementation of convolution filters if $\beta_N(\lambda) < 1$, and
- frequency implementation of multiplicative filters if $\beta_N(\lambda) > 1$.

Since $\beta_N(\lambda)$ is a monotonically decreasing function of λ given fixed N, the trade-off scale at which global frequency and temporal aliasing errors are balanced equals

$$\lambda_* = \frac{1}{2}\log\frac{N}{2\pi}\,. \tag{28}$$

To see this, note that the function $\|\bar{\varepsilon}(\,.\,;\xi)\|$ appearing on the r.h.s. of Eq. (27) is monotonic, whereas ρ and η are inversely proportional. Numerator and denominator are therefore balanced iff $\rho = \eta$, in other words, using $\eta = 2\pi/(N\rho)$, iff $\rho^2 = 2\pi/N$. Using $\rho = T/\sigma$, the result for logarithmic scale then follows from Eq. (22).

We may conclude that in the fine scale region $\lambda < \lambda_*$ the frequency implementation is the preferred one, whereas for coarse scales $\lambda > \lambda_*$ one should consider the spatial implementation. Thus the advantage of scale independent computation time of the Fourier method comes at the price of suboptimal accuracy, at least for the upper part of the scale interval. Because it is a balance of logarithmic scales, the trade-off scale appears rather small, *e.g.* for a signal consisting of $N = 256$ samples we have $\sigma_* \approx 6.4\,T$.

3 Conclusion

We have investigated accuracy trade-offs between spatial and Fourier implementations of 1D linear filters. We have focused on the Gaussian scale-space paradigm, for which explicit results have been stated. In particular, we have derived a simple formula for the trade-off scale, $\lambda_* = \frac{1}{2}\log\frac{N}{2\pi}$, which divides the logarithmic scale interval into two subintervals. Scales in the lower part of this interval, $\lambda < \lambda_*$, are most accurately extracted by the Fourier method, whereas scales in the upper part, $\lambda > \lambda_*$, require spatial filtering. Thus in computing a scale-space representation one should always balance the benefit of quality against the concomitant cost of computational load imposed by the requirement of large convolution filters.

References

[1] T. Lindeberg, "Scale-space for discrete signals," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 3, pp. 234–245, 1990.

[2] T. Lindeberg, *Scale-Space Theory in Computer Vision*. The Kluwer International Series in Engineering and Computer Science, Kluwer Academic Publishers, 1994.

[3] R. Deriche, "Fast algorithms for low-level vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 1, pp. 78–87, 1990.

[4] R. Deriche, "Recursively implementing the Gaussian and its derivatives," in *Proceedings of the 2nd Singapore International Conference on Image Processing (Singapore, September 1992)* (V. Srinivasan, O. Sim Heng, and A. Yew Hock, eds.), pp. 263–267, Singapore: World Scientific, 1992.

[5] W. T. Cochran, J. W. Cooley, D. L. Favin, H. D. Helms, R. A. Kaenel, W. W. Lang, G. C. Maling, D. E. Nelson, C. M. Rader, and P. D. Welch, "What is the Fast Fourier Transform?," in *Proceedings of the IEEE*, pp. 1664–1674, October 1967.

[6] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C; the Art of Scientific Computing*. Cambridge: Cambridge University Press, 1988.

[7] L. M. J. Florack, *Image Structure*, vol. 10 of *Computational Imaging and Vision Series*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1997.

[8] B. M. t. Haar Romeny, L. M. J. Florack, J. J. Koenderink, and M. A. Viergever, eds., *Scale-Space Theory in Computer Vision: Proceedings of the First International Conference, Scale-Space'97, Utrecht, The Netherlands*, vol. 1252 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag, July 1997.

[9] J. Sporring, M. Nielsen, L. M. J. Florack, and P. Johansen, eds., *Gaussian Scale-Space Theory*, vol. 8 of *Computational Imaging and Vision Series*. Dordrecht: Kluwer Academic Publishers, 1997.