Theoretical Foundations of Gaussian Convolution by Extended Box Filtering

Pascal Gwosdek[1], Sven Grewenig[1], Andrés Bruhn[2], and Joachim Weickert[1]

[1] Mathematical Image Analysis Group, Dept. of Mathematics and Computer Science,
Campus E1.1, Saarland University, 66041 Saarbrücken, Germany
{gwosdek,grewenig,weickert}@mia.uni-saarland.de
[2] Vision and Image Processing Group,
Cluster of Excellence Multimodal Computing and Interaction,
Saarland University, Campus E 1.1, 66041 Saarbrücken, Germany
bruhn@mmci.uni-saarland.de

Abstract. Gaussian convolution is of fundamental importance in linear scale-space theory and in numerous applications. We introduce iterated extended box filtering as an efficient and highly accurate way to compute Gaussian convolution. Extended box filtering approximates a continuous box filter of arbitrary non-integer standard deviation. It provides a much better approximation to Gaussian convolution than conventional iterated box filtering. Moreover, it retains the efficiency benefits of iterated box filtering where the runtime is a linear function of the image size and does not depend on the standard deviation of the Gaussian. In a detailed mathematical analysis, we establish the fundamental properties of our approach and deduce its error bounds. An experimental evaluation shows the advantages of our method over classical implementations of Gaussian convolution in the spatial and the Fourier domain.

Keywords: Gaussian scale-space, box filter, image processing, computer vision.

1 Introduction

Convolution with a Gaussian is one of the most widely used linear filter operations in signal and image processing. It forms the backbone of Gaussian scale-space theory [4,9,13] which has been introduced in Japanese and English papers of Iijima [8] long before it became popular in the western world by Witkin's work [16]. The strong regularisation properties of Gaussian convolution render the filtered signal infinitely times differentiable and stabilise the numerical evaluation of higher order derivatives. Gaussian convolution is inevitable for the detection of edges [2,11] and interest points [6,10] that play a central role in computer vision. The rapid decay properties of the Gaussian both in the spatial and the Fourier domain and the fact that it is the only filter that is rotationally invariant and separable under convolution make Gaussian convolution a perfect low-pass filter in linear systems theory.

A.M. Bruckstein et al. (Eds.): SSVM 2011, LNCS 6667, pp. 447–458, 2012.

Many applications require an accurate and efficient implementation of Gaussian convolution in order to ensure the high quality of the results, to meet runtime requirements, or even to guarantee convergence. However, this can be challenging: It comes down to a convolution of the input signal with a kernel function with infinite support. The m-dimensional Gaussian kernel

$$K_\sigma(\boldsymbol{x}) \;=\; \frac{1}{(2\pi\sigma^2)^{\frac{m}{2}}} \;\exp\left(-\frac{|\boldsymbol{x}|^2}{2\cdot\sigma^2}\right) \tag{1}$$

of standard deviation σ has a characteristic 'bell curve' shape which drops off rapidly towards $\pm\infty$. This is why in practice one often applies a discrete convolution with a sampled and renormalised kernel that is cut off at $n \cdot \sigma$. However, this method becomes inefficient for large σ, as the number of operations grows linearly in the number of samples of both the signal and the kernel. A more efficient alternative for those cases is the computation as a point-wise multiplication in the frequency domain [1]. To this end, a Fourier transform is applied to both the kernel and the signal, the multiplication is performed, and the result is transformed back into the spatial domain. Since the Gaussian kernel in the frequency domain can immediately be evaluated, this method reduces to two fast Fourier transforms, and one point-wise multiplication.

Although these spatial and Fourier-based implementations are the most popular algorithms for Gaussian convolution, and their trade-offs are well investigated [5], there are also further alternatives: Approximations with recursive filters [3,17] offer a runtime behaviour that scales linearly in the number of pixels. However, these filters require a special boundary treatment and a higher implementational effort than other methods which poses additional challenges [14]. Since Gaussian scale-space is equivalent to evolving the image under a homogeneous diffusion problem, one can also implement Gaussian convolution with efficient numerical methods for partial differential equations, e.g. with implicit finite difference schemes [7]. Unfortunately, this requires the fast solution of linear systems of equations which is also a nontrivial task. Gaussian convolution can also be approximated by discrete convolution with binomial kernels. They have a finite support and offer some interesting properties from an implementational viewpoint, but do not allow to approximate Gaussians with arbitrary standard deviations. This can constitute a drawback in scale-space applications which aim at representations at arbitrary scales.

A simple but extremely fast discrete approximation of Gaussian smoothing can be achieved by convolution with iterated box filters [15]. A box filter uses a normalised kernel with identical coefficients within its finite support. By the central limit theorem, a sufficiently high number of iterations with a box filter approximates a Gaussian arbitrarily well. However, this has the same drawback as convolution with binomial kernels: It introduces a quantisation to the range of standard deviations that can be approximated.

In our paper we address this problem. We advocate a modification of the box filter that is based on a new discretisation of the continuous box kernel. In particular, we concentrate on establishing important properties of the resulting *extended box filter*: It combines the simplicity and algorithmic efficiency of the