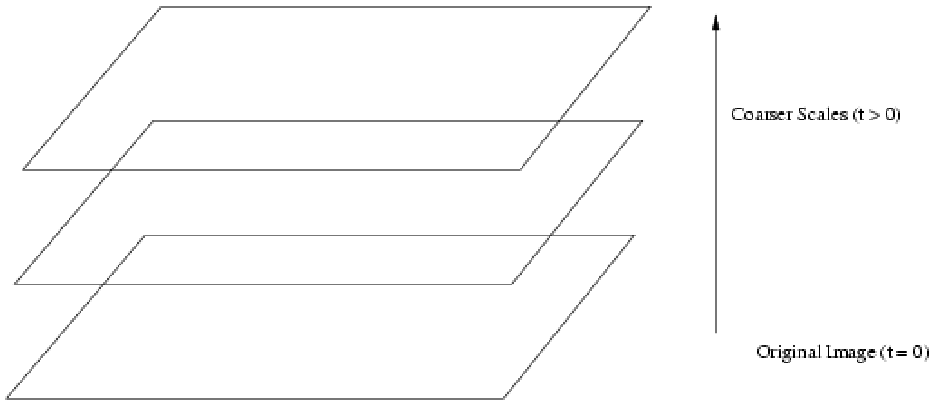

Figure 1: Fine To Coarse Scales. The parameter t is often associated with the generation of a scale-space family, where $t = 0$ denotes the original image and $t > 0$ increasing (coarser) scales.

Scale-Space Theory Notes

This document is largely a summary of ideas taken from Tony Lindeberg's paper: "Scale-Space for Discrete Signals" and Hummel's paper: "Representations based on Zero-crossings in Scale-Space". The essential idea in scale-space theory is to systematically remove details from an image so as to describe its structure in a fine-to-coarse (multi-resolution) fashion. Another way of thinking about this is to obtaiin an ordered set of derived images intended to represent the original at increasing (coarser) levels of scale, as illustrated in Figure 1.

1 1D Gaussian Blurring

Let $f : \mathcal{R} \rightarrow \mathcal{R}$ be a continuous signal. Let $L : \mathcal{R} \times \mathcal{R}_+ \rightarrow \mathcal{R}$ be a function defined by convolution with a Gaussian kernel: $g : \mathcal{R} \times \mathcal{R}_+/\{0\} \rightarrow \mathcal{R}$ such that

$$L(x; 0) = f(x)$$

$$L(x;t) = \int_{-\infty}^{\infty} (1/\sqrt{2\pi t}) e^{-\zeta^2/2t} f(x - \zeta) d\zeta \quad \text{if} \quad t > 0.$$

This family is equivalent to that obtained by the heat or diffusion equation:

$$L(x; 0) = f(x)$$

$$\frac{\partial L}{\partial t} = \frac{1 \partial^2 L}{2 \partial x^2}$$

The properties of this process can be enumerated as follows:

1. As t increases no new local extrema (or zero-crossings) appear. This is referred to as "structure simplification".

2. $L(x; t_2)$ depends exclusively on $L(x; t_1)$ if $t_2 > t_1 \geq 0$. This is a form of the causality principle.

3. The result of this convolution (in the Gaussian smoothing view) is shift-invariant and holds for all $f(x)$. The 2D version leads to the familiar diffusion equation, which we studied in class, and which is also the subject of the paper by Hummel.

2 Scale-Space Axioms

The goal of Lindeberg's paper is to develop a theory for scale-space which applies to discrete signals, rather than the classical theory which formulated in the continuous domain. He begins with the following axiom

Axiom 1 *A 1D discrete kernel $K : Z \to \mathcal{R}$ is denoted a scale-space kernel if for all signals $f_{\text{in}} : Z \to R$ the number of local extrema in $f_{\text{out}} = K * f_{\text{in}}$ is non-increasing.*

The basic assumptions are:

1. All signals are real-valued functions defined on the same (infinite) grid. These functions take the form $f : Z \to \mathcal{R}$.
2. All representations should be generated by convolution of the original image with a kernel such that the assumptions of linearity and shift invariance hold.
3. $t = 0$ represents the original signal and $t \uparrow$ corresponds to coarser scale.

We have the following propostion

Proposition 1 *Let $K : Z \to \mathcal{R}$ be a scale-space kernel and $\mathcal{L}$ be a linear operator from the space of real-valued discrete functions to itself. Then for any $f : Z \to R$ the number of local extrema in $\mathcal{L}(K * f)$ cannot exceed the number of local extrema in $\mathcal{L}(f)$.*

This proposition follows because K and $\mathcal{L}$ commute, i.e.,

$$\mathcal{L}(K * f) = K * \mathcal{L}(f),$$

where $K * \mathcal{L}(f)$ has fewer zero-crossings than $\mathcal{L}(f)$.

A special case of such a linear operator is the operator $g = f(x+1) - f(x)$ (a first-order discrete derivative) whose zero-crossings correspond to local extrema of $f(x)$. This implies that the number of zero-crossings of $f(x)$ is non-increasing. In fact a function as well as all its derivatives will become "smoother" in the sense that no new local extrema will be created.

3 Qualitative Properties of Discrete Scale-Space Kernels

Consider a test function $f_n(x)$ defined as a discrete delta function:

$$f_{in}(x) = \delta(x) = \{1, \quad \text{if} \quad x = 0; = 0 \quad \text{otherwise}$$

We know that convolution of a signal by a delta function leaves the signal unchanged. In other words, for a discrete scale space kernel K

$$(K * \delta)(x) = K(x).$$

But since a scale-space kernel must admit any input function and since the delta function has only one local maximum and no zero-crossings, $K(x)$ must have at most one extremum and no zero-crossings. Therefore we have:

Proposition 2 *All coefficients of a discrete scale-space kernel must have the same sign. Furthermore, the coefficient sequence $\{K(n)\}_{-\infty}^{\infty}$ must be unimodal.*

It is also useful to impose the requirements that $K \in l_1$, i.e.,

$$\sum_{n=-\infty}^{\infty} |K(n)|$$

is finite and that f_{in} is bounded. These requirements allow one to speak of the fourier transform of the kernel coefficient sequence and hence to determine properties of scale-space kernels in the frequency domain.

3.1 Generalized Binomial Kernels

Such kernels are now considered and are defined as follows:

$$K^{(2)}(n) = \{p, n = 0; q, n = -1; 0 \quad \text{otherwise}\}.$$

We assume that $p, q \geq 0$ and that $p + q = 1$, without loss of generality. It is easy to see that convolutioon of a test function f_{in} with a kernel of the form $K^(2)$ gives

$$\frac{f_{in}(x) - f_{in}(x-1)}{f_{out}(x) - f_{in}(x-1)} = (p+q)/q$$

or geometrically, the result in Figure 3 of Lindeberg's paper. It should be clear that no new zero-crossings can be introduced by sucha a transformation. Hence a kernel of the type $K^{(2)}$ is a discrete scale-space kernel. Furthermore, the following is also true

Proposition 3 *All kernels obtained by repeated convolution of 2-kernels by 2-kernels are scale-space kernels denoted $*_{i=1}^{n} K_i^{(2)}$.*

Furthermore, there is a generating function form of such kernels which is elaborated on in proposition 5 of Lindeberg's paper.

3.2 Properties of Kernel Coefficients

Let $f_{out} = K * f_{in}$ be represented in matrix form as $f_{out} = C f_{in}$ where C is a $N \times N$ matrix with values $K(i-j)$ along its diagonals (it is a Toeplitz matrix). Then we have the following propositions, the proofs of which are spelled out in Lindeberg's paper:

Proposition 4 *The eigenvalues of C are non-negative.*

Proposition 5 *The fourier transform $\Psi_k(\theta) = \sum_{n=-\infty}^{\infty} K(n)e^{-in\theta}$ of a symmetric discrete scale-space kernel with finite support is real-valued and is non-negative.*

Proposition 6 *The fourier transform of a symmetric discrete scale-space kernel with finite support is unimodal on the interval $[-\pi, \pi]$ with maximum value at $\theta = 0$.*

4 Kernel Classification

This section derives an explicit characterization of discrete scale-space kernels. The main result is that the generalized binomial kernels introduced earlier are the only discrete scale-space kernels with finite support.

Theorem 1 *A discrete scale-space kernel $K : Z \rightarrow \mathcal{R}$ is a scale-space kernel IFF all minors of the matrix C with entries $K(i-j)$ along its diagonals are non-negative. The minor of a matrix is the reduced determinant of a determinant expansion, i.e., the determinant of the matrix obtained by deleting the row and column which intersect a particular element. The sequence of such filter coefficients $\{K(n)\}_{n=-\infty}^{n=\infty}$ is called a normalized Polya frequency sequence.*

Theorem 2 *The second theorem gives an explicit characterization of a normalized Polya frequency sequence in terms of its generating function $\phi_K(z) = \sum_{n=-\infty}^{\infty} K(n)z^n$, which has a form described by Eq. 16 in the paper.*

Theorem 3 *The kernels with the form $*_{i=1}^{n} K(n)$ are (except for rescaling and translation) the only discrete-scale space kernels wt finite support.*

As a corollary, convolution with any finite scale-space kernel can be decomposed into convoution with kernels having two strictly positive consecutive filter coefficients.

5 Axiomatic Scale-Space Construction

A scale-space kernel with discrete scale parameter t results in the process

$$L_0 = f \quad \text{thefinestscale}$$

$$L_i = L_{i-1} * K_i$$

where the kernels are discrete scale-space kernels selected for suitable amounts of "blurring". However, a key issue is how to select the discrete scales. Rather than adopt this approach we choose to treat t as a continuous scale parameter (i.e., to process all scales). In other words, we postulate that scale-space is generated by a 1 parameter family of kernels, i.e.,

$$L(x; 0) = f(x)$$

$$L(x;t) = \sum_{n=-\infty}^{\infty} T(n;t) f(x-n) \quad t > 0 \quad \text{(adiscreteconvolution)}$$

This formulation reflects the linear shift-invariance assumption and in addition $T(;t)$ is assumed to be a scale-space kernel. In order to simplify the analysis, we impose a semi-group requirement on the family of kernels, i.e.,

$$T(.;s) * T(.;t) = T(.;s+t).$$

Note that if $t_2 > t_1$, this requirement implies that

$$\begin{aligned} L(.;t_2) &= T(.;T_2) * f \\ &= (T(.;t_2 - t_1) * T(.;t_1)) * f \\ &= T(.;t_2 - t_1) * (T(.;t_1) * f) \\ &= T(.;t_2 - t_1) * L(.;t_1). \end{aligned}$$

Thus the representation at a coarse scale can be calculated from the representation at any finer scale.

The maiin theorem, which is stated in its full glory as Theorem 4 in the paper, is that the above conditions combined with a normalization criterion:

$$\sum_{n=-\infty}^{\infty} T(n[t) = 1$$

and a symmetry constraint

$$T(-n;t) = T(n;t)$$

determine the family of kernels upto a scaling parameter α:

$$T(n;t) = e^{-}\alpha t I_n(\alpha t)$$

where

$$I_n(t) = I_{-n}(t) = (-1)^n J_n(it); n \geq 0, t > 0.$$

Here the I_n's are modified Bessel functions of integer order and $J_n(it)$'s are Bessel functions with purely imaginary arguments. (See mathworld.wolfram.com for a discussion of these functions). Unfortunately these functions can be described by generatiing functions, but they cannot be easily expressed in closed form. Fortunately there is a recurrence relation which can be exploited which can lead to a form of a diffusion equation, without explicitly having to use the Bessel functions explicitly. This is explained in the next section.

6 Numerical Implementation for the 1D Case

Starting with $L(x;0) = f(x)$ we need to derive

$$L(x;t) = \sum_{n=-\infty}^{\infty} T(n;t) f(x-n) \quad t > 0,$$

where

$$T(n;t) = e^{-\alpha t} I_n(\alpha t).$$

One way to go now is to exploit the property that the filter coefficients satisfy the recurrence relation:

$$I_{n-1}(t) - I_{n+1}(t) = (2n/t) I_n(t).$$

A reasonable approach is to truncate the infinite sum for sum sufficiently large value of index N, and to estimate filter coefficients as one needs them using the above recurrence relation, as described in some detail on p. 242. However, one might choose also to solve the problem by showing that it is equivalent to a form of a discrete diffusion equation.

More formally, for modified Bessel functions we have the relation

$$2I'_n(t) = I_{n-1}(t) + I_{n+1}(t).$$

It is easy to show that with the kernel $T(n;t) = e^{-t} I_n(t)$

$$\begin{aligned} \frac{\partial T(n;t)}{\partial t} &= -e^{-t} I_n(t) + e^{-t} I'_n(t) \\ &= e^{-t}(1/2 I_{n-1}(t) + 1/2 I_{n+1}(t) - I_n(t)) \\ &= 1/2(T(n-1;t) + T(n+1;t) - 2T(n;t)) \qquad (1) \end{aligned}$$

Because the kernel satisfies the above system of equations it is easy to show that with initial conditions

$$L(x;0) = f(x)$$

$$\frac{\partial L(x;t)}{\partial t} = 1/2((L(x+1;t) - 2L(x;t) + L(x-1;t)) \quad x \in Z$$

which is Theorem 6 in the paper. This equation can be numerically discretized and solved in a straightforward fashion.

The remainder of Lindeberg's paper discusses the extension from one to two dimensions. I will expect that you have read this section to appreciate at least the subtleties involved.