

Laplacian Eigenimages in Discrete Scale Space

Martin Tschirsich¹ and Arjan Kuijper^{1,2}

¹ Technische Universität Darmstadt, Germany

² Fraunhofer IGD, Darmstadt, Germany

Abstract. Linear or Gaussian scale space is a well known multi-scale representation for continuous signals. However, implementational issues arise, caused by discretization and quantization errors. In order to develop more robust scale space based algorithms, the discrete nature of computer processed signals has to be taken into account. Aiming at a computationally practicable implementation of the discrete scale space framework we used suitable neighborhoods, boundary conditions and sampling methods. In analogy to prevalent approaches, a discretized diffusion equation is derived from the continuous scale space axioms adapted to discrete two-dimensional images or signals, including requirements imposed by the chosen neighborhood and boundary condition. The resulting discrete scale space respects important topological invariants such as the Euler number, a key criterion for the successful implementation of algorithms operating on its deep structure. In this paper, relevant and promising properties of the discrete diffusion equation and the eigenvalue decomposition of its Laplacian kernel are discussed and a fast and robust sampling method is proposed. One of the properties leads to Laplacian eigenimages in scale space: Taking a reduced set of images can be considered as a way of applying a discrete Gaussian scale space.

1 Introduction

In the field of computer vision, deriving information from observed images or *signals* is a central problem. Various strategies have been invented to do so in a performant manner, usually by applying some kind of operator. These operators often detect or rely on the presence of image structure or features such as edges or stationary points and are of fixed size. Their performance then depends on the *inner scale*, the sampling density or resolution of the image they operate on.

To overcome this dependence between operator and inner scale, various strategies such as integral images used by the Viola-Jones Framework [1] or a whole range of multi-scale representations have been proposed [2]. Almost all those strategies consist of transforming the given images into a scale independent representation first before applying an operator on this representation. A common requirement for such preliminary transformations is to mask as little information present in the original image as possible, that is, not to rely on prior information not present in the image data itself and therefore not unnecessarily limiting the range of tasks they can be applied to.

Various scale-invariant or multi-scale signal representations satisfying these requirements exist, such as the lowpass pyramid, wavelet and scale space representations, however there are qualitative differences. A pyramid representation consists of several fixed

images or *scales* with preselected resolution, each computed by smoothing and subsampling a finer scale with the first and finest scale being the initial image. A sliding window approach is then used to evaluate a fixed-size operator at every position and on every scale of the pyramid. Because of its fixed number and decreasing size of scales, the pyramid representation can be efficiently computed and stored. The subsampling operator however, primarily implemented for performance reasons, introduces often undesired subsampling artifacts. It also complicates and sometimes prevents tracing of image features over multiple scales. The Gaussian scale space is used to overcome these drawbacks and can be thought of as the natural generalization of the lowpass pyramid. It is also popular for its theoretical foundation. A Gaussian scale space representation of a given signal is a family of derived signals, progressively smoothed with a Gaussian filter.

The focus of this paper is on the linear or Gaussian scale space representation of discrete images or signals and the various implementational issues that have to be addressed in order for this representation to be useful and applicable to real world data.

The discrete scale space proposed by Lindeberg [3] takes the discrete nature of computer processed signals into account. It is based on equivalent assumptions and axioms that have been used to derive the continuous Gaussian scale space adapted to discrete signals. It is our belief that porting scale space algorithms from a discretized continuous to the discrete scale space will eventually lead to more accurate, robust and possibly faster implementations. The discrete scale space formalized by Lindeberg however does not respect important topological invariants such as the Euler number [4]. Since most algorithms that operate on the deep structure of the Gaussian scale space require this topological invariant to hold [5, 6], we had to give a modified definition of the discrete scale space respecting the Euler number [7]. In this paper we investigate and discuss relevant and promising properties of the discrete diffusion equation and the eigenvalue decomposition of its Laplacian kernel. We propose a fast and robust sampling method. One of the properties leads to what we coined as Laplacian eigenimages in scale space, where taking a reduced set of images can be considered as a way of applying a discrete Gaussian scale space.

2 Continuous and Discrete Scale Space

Linear or *Gaussian scale space* was introduced in western literature by Witkin [8] and extended into two dimensions by Koenderink [9] and has become a useful framework for multi-scale signal representation. However, Gaussian scale space was first described in Japan as of 1959 [10]. In 1962, Taizo Iijima proposed convolution with the Gaussian kernel as the canonical way to construct the Gaussian scale space [11]. Other, more general scale spaces exist besides the Gaussian scale space. Hereafter, scale space will refer to the two-dimensional linear Gaussian scale space.

2.1 Continuous Scale Space

The *scale space representation* of a continuous two-dimensional signal $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a one-parameter family of derived signals $L : \mathbb{R}^2 \times \mathbb{R}_+$ defined by one of the following equivalent definitions.

Convolution with a Scale Space Kernel: $L(\cdot, \cdot; t) = g(\cdot, \cdot; t) * f(\cdot, \cdot)$ where $g(x, y; t) = \frac{1}{4\pi t} e^{-\frac{x^2+y^2}{4t}}$. Commonly, the factor $4t$ is taken instead of $2\sigma^2$. The Gaussian g is also called the canonical scale space kernel.

Scale Space Axioms: A common definition of the scale space is by a limited set of scale space axioms derived from real world requirements. Koenderink [9] formulated the axioms of causality, homogeneity and isotropy. Another equivalent set of axioms states that $L(\cdot, \cdot, t) = \mathcal{L}_t f(\cdot, \cdot)$ where \mathcal{L}_t is a linear shift invariant operator and thus representable as a convolution. Together with non-enhancement of local extrema stating that $\partial_t L < 0$ for maxima and $\partial_t L > 0$ for minima and some other additional requirements, these axioms all lead to the Gaussian kernel as a unique choice as the scale space kernel. An overview of a wide range of continuous Gaussian scale-space axiomatics used by different authors can be found in [10].

Diffusion Equation: The scale space representation L of a signal f is the solution of the diffusion equation

$$\partial_t L = \Delta L = \partial_{xx} L + \partial_{yy} L$$

with $L(\cdot; 0) = f(\cdot)$ as initial condition. Δ denotes the Laplacian operator.

2.2 Discrete Scale Space

In this section, we will give a concise definition of the Gaussian scale space for two-dimensional discrete signals comprising a continuous scale parameter. The scale space representation of a discrete signal $f : \mathbb{Z}^2 \rightarrow \mathbb{R}$ is a one-parameter family of derived signals $L : \mathbb{Z}^2 \times \mathbb{R}^+$ defined by one of the following equivalent definitions:

Convolution with a Scale Space Kernel: $L(\cdot, \cdot; t) = k(\cdot, \cdot; t) * f(\cdot, \cdot)$ where $k : \mathbb{Z}^2 \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is the discrete version of the Gaussian scale space kernel.

$$L(x, y; t) = \sum_{m=-\infty}^{\infty} k(m; t) \sum_{n=-\infty}^{\infty} k(n; t) f(x - m, y - n)$$

Scale Space Axioms: For a complete set of axioms describing the discrete scale space, we refer to those chosen by Lindeberg [3].

Diffusion Equation: The scale space representation L of a signal f is the solution of the diffusion equation

$$\partial_t L = \alpha \nabla_2^5 L + \beta \nabla_{\times}^2 L \tag{1}$$

with $L(\cdot, \cdot; 0) = f(\cdot, \cdot)$ as initial condition and some $\alpha, \beta \geq 0$. The five-point operator ∇_2^5 and the cross operator ∇_{\times}^2 approximate the continuous Laplacian operator. They

correspond to convolution of $L(\cdot, \cdot; t)$ with kernel $\begin{bmatrix} 1 & & \\ 1 & -4 & 1 \\ 1 & & \end{bmatrix}$ respectively $\begin{bmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & -2 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}$

for fixed scale t . For $\alpha = 1$ and $\beta = 0$ (1) simplifies to $\partial_t L = \nabla_2^5 L$ and results in k being separable.

3 Eigenbasis Decomposition of the Laplacian

The discrete scale space representation L of signal f holds various useful properties. Sampling of the scale space for example requires the ability to efficiently compute a scale $L(\cdot, \cdot; t_0)$ for fixed t_0 and relies on certain characteristics of the Laplacian operator.

The discrete scale space representation L of f is continuous in scale t . A computational investigation of L however must rely on a finite number of sampled scales. There are multiple approaches to sampling L differing in accuracy, runtime complexity and memory usage. One apparent approach is given by the definition of L via discrete convolution with a scale space kernel. The scale space kernel is of infinite domain and must be truncated in order to compute an individual scale, thus introducing truncation errors. A periodic boundary condition for f further complicates the computation. In this case, circular convolution with a Laplacian kernel provides for a more elegant but still computationally complex solution. Applied in its eigenspace however, the circular convolution operator reduces to a simple and much less complex scaling transformation. This section details how to efficiently decompose a scale of L and its derivative $\partial_t L$ into a sum of eigenimages of the Laplacian circular convolution operator. The next section then provides a simple solution of the discretized diffusion equation, enabling for fast and accurate sampling of L .

For periodic discrete signals f with discrete domain $D(f) = [1, M] \times [1, N]$, the diffusion equation $\partial_t L = \nabla_5^2 L$ can be written as a circular convolution with finite Laplacian kernel

$$\partial_t L = \nabla_5^2 L = \begin{bmatrix} 1 & & \\ 1 & -4 & 1 \\ & & 1 \end{bmatrix} \circledast L,$$

where \circledast denotes the circular convolution operator.

The discrete circular convolution is a linear operator and can be expressed in matrix form if we consider $L(\cdot, \cdot; t)$ to designate a vector. Scale $L(\cdot, \cdot; t)$ of the scale space representation L can be represented as a vector $\mathbf{L}(t) \in \mathbb{R}^{MN}$ with $\mathbf{f} = \mathbf{L}(0)$.

$$\mathbf{L}(t) = \begin{bmatrix} L(1, 1; t) \\ L(1, 2; t) \\ \vdots \\ L(M, N; t) \end{bmatrix} \in \mathbb{R}^{MN}$$

For periodic f , the diffusion equation can be written in matrix form as $\partial_t L = \nabla_5^2 L = \Delta_{M,N} \mathbf{L}$ where $\Delta_{M,N} \in \mathbb{R}^{MN \times MN}$ denotes a circulant block matrix corresponding to the Laplacian operator ∇_5^2 . For $M, N \geq 3$ it takes the form

$$\Delta_{M,N} = \begin{bmatrix} \mathbf{A}_N & \mathbf{I}_N & & \mathbf{I}_N \\ \mathbf{I}_N & \ddots & \ddots & \\ & \ddots & \ddots & \mathbf{I}_N \\ \mathbf{I}_N & & \mathbf{I}_N & \mathbf{A}_N \end{bmatrix} \in \mathbb{R}^{MN \times MN}, \mathbf{A}_N = \begin{bmatrix} -4 & 1 & & 1 \\ 1 & \ddots & \ddots & \\ & \ddots & \ddots & 1 \\ 1 & & 1 & -4 \end{bmatrix} \in \mathbb{R}^{N \times N},$$

where \mathbf{I}_N is the identity of $\mathbb{R}^{N \times N}$. Let the normalized eigenvectors and eigenvalues of Δ_{MN} be $\mathbf{u}_{i,j}$ and $\lambda_{i,j}$. Since Δ_{MN} is a real, symmetric and therefore diagonalizable matrix, its $M \cdot N$ orthonormal eigenvectors $\mathbf{u}_{i,j}$ form an eigenbasis $\mathbf{U} = [\mathbf{u}_{1,1}, \mathbf{u}_{2,1}, \dots, \mathbf{u}_{M,N}]$ and $\Delta_{MN} = \mathbf{U}\Lambda\mathbf{U}^T$ with $\Lambda = \text{diag}(\lambda_{1,1}, \lambda_{2,1}, \dots, \lambda_{M,N})$ and $\lambda_{i,j} \leq 0$ since Δ_{MN} is a negative-semidefinite matrix.

The scale $\mathbf{L}(t)$ of the scale space representation \mathbf{L} can be written as a weighted sum of eigenimages of the Laplacian operator, i.e. as a scalar product of the orthonormal eigenvectors $\mathbf{u}_{i,j}$ of $\Delta_{M,N}$ and the scalar coefficients $c_{i,j}(t) = \langle \mathbf{L}(t), \mathbf{u}_{i,j} \rangle$ resulting from the projection of $\mathbf{L}(t)$ to $\mathbf{u}_{i,j}$:

$$\mathbf{L}(t) = \sum_{i,j} c_{i,j}(t) \mathbf{u}_{i,j}$$

Its partial derivative $\partial_t \mathbf{L}(t)$ can then be computed from scaling each projected component separately by the corresponding eigenvalue.

$$\partial_t \mathbf{L}(t) = \mathbf{U}\Lambda\mathbf{U}^T \mathbf{L}(t) = \sum_{i,j} c_{i,j}(t) \lambda_{i,j} \mathbf{u}_{i,j}.$$




Fig. 1. Eigenimages and eigenvalues of the Laplacian $\Delta_{MN} \in \mathbb{R}^{6 \times 6}$ and $\Delta_{MN} \in \mathbb{R}^{10 \times 10}$

4 Efficient Computation of the Eigenimages

The size and values of the Laplacian matrix $\Delta_{M,N}$ depend uniquely on the dimension of f equal to the cardinality of $D(f)$. They are easily computed following the specification given in the previous section. Unfortunately, even for images of moderate size of e.g. 128×128 pixel, f would have $2^{7 \cdot 2}$ dimensions, resulting in a Laplacian matrix of 2^{28} entries taking up to $2GB$ memory. Even under consideration of less computationally and memory intensive sparse matrix formats, such matrices are too large to be

handled efficiently. Therefore, in order to reduce the overall complexity of e.g. computing the eigenbasis \mathbf{U} , we have to exploit the symmetric structure of the Laplacian. The Laplacian kernel is of rank > 1 , thus not separable. However, using the distributivity of the circular convolution, it can be rewritten as a summation of two one dimensional kernels.

$$\partial_t L = \begin{bmatrix} 1 \\ 1 & -4 & 1 \\ 1 \end{bmatrix} \circledast L = \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix} \circledast L + [1 \ -2 \ 1] \circledast L$$

In matrix form, this translates to a direct summation of two substantially smaller matrices. The Laplacian matrix $\Delta_{M,N} \in \mathbb{R}^{MN \times MN}$ can be written as the direct sum of two ∇_3^2 operators $\Delta_M \in \mathbb{R}^{M \times M}$ and $\Delta_N \in \mathbb{R}^{N \times N}$

$$\Delta_{M,N} = \Delta_M \oplus \Delta_N = (\Delta_M \otimes \mathbf{I}_N) + (\mathbf{I}_M \otimes \Delta_N)$$

where Δ_M and Δ_N are the matrix representations of the row wise applied central difference operator of second order. They differ only in their dimensions. \otimes denotes the Kronecker product. For $M \geq 3$, Δ_M has the form of a Toeplitz matrix.

$$\Delta_M = \begin{bmatrix} -2 & 1 & & 1 \\ 1 & -2 & 1 & \\ & \ddots & \ddots & \ddots \\ 1 & & 1 & -2 \end{bmatrix} \in \mathbb{R}^{M \times M}$$

Each eigenvector $\mathbf{u}_{i,j}$ of $\Delta_{M,N}$ can be expressed as the outer product of two eigenvectors \mathbf{v}_i and \mathbf{w}_j of Δ_M and Δ_N . The corresponding eigenvalue $\lambda_{i,j}$ is then the sum of the corresponding eigenvalues v_i and ω_j of Δ_M and Δ_N .

$$\begin{aligned} \Delta_{M,N} \mathbf{u}_{i,j} &= \lambda_{i,j} \mathbf{u}_{i,j} \\ \Leftrightarrow (\Delta_M \oplus \Delta_N) (\mathbf{v}_i \otimes \mathbf{w}_j) &= (v_i + \omega_j) (\mathbf{v}_i \otimes \mathbf{w}_j) \end{aligned}$$

Omitting all details that can be found in [12], we finally have an analytic formula expressing the eigenvalues $\lambda_{i,j}$ and eigenvectors $\mathbf{u}_{i,j}$ of the Laplacian matrix $\Delta_{M,N}$ with $i = 0 \dots M - 1, j = 0 \dots N - 1$.

$$\lambda_{i,j} = v_i + \omega_j = \frac{(\lambda_{M,i}^F)^2}{(\lambda_{M,i}^F + 1)} + \frac{(\lambda_{N,j}^F)^2}{(\lambda_{N,j}^F + 1)} = \frac{(e^{(\frac{2\pi\ell}{M}i)} - 1)^2}{e^{(\frac{2\pi\ell}{M}i)}} + \frac{(e^{(\frac{2\pi\ell}{N}j)} - 1)^2}{e^{(\frac{2\pi\ell}{N}j)}}$$

$$\begin{aligned} \mathbf{u}_{i,j} &= v_i \otimes w_j = d_{M,i}^F \otimes d_{N,j}^F = d_{M,i}^F = \\ &\left[\exp\left(\frac{2\pi\ell}{M}i\right)^0, \dots, \exp\left(\frac{2\pi\ell}{M}i\right)^{M-1} \right] \otimes \left[\exp\left(\frac{2\pi\ell}{N}j\right)^0, \dots, \exp\left(\frac{2\pi\ell}{N}j\right)^{N-1} \right] \end{aligned}$$

These eigenvectors are not guaranteed to be real, although $\Delta_{M,N}$ always possesses a real eigenbase. Exploiting further symmetries of the eigenvectors v_i and w_j , we can derive a real eigenbase $\tilde{\mathbf{U}}$ given by $\tilde{\mathbf{u}}_{i,j} = \tilde{v}_i \otimes \tilde{w}_j$ with $\tilde{v}_i = \Re(v_i) + \Im(v_i)$ and $\tilde{w}_j = \Re(w_j) + \Im(w_j)$.

5 Discrete Scale Space Eigenbasis Decomposition

As shown in a previous section, scale $\mathbf{L}(t)$ and its partial derivative $\partial_t \mathbf{L}(t)$ of the scale space representation \mathbf{L} can be written as a weighted sum of the eigenimages $\mathbf{u}_{1,1}, \dots, \mathbf{u}_{M,N}$ of the Laplacian operator. This implicit change of base allows us to give a simple solution for the discretized diffusion equation.

$$\begin{aligned} \partial_t L(t) &= \Delta_5^2 L(t) \\ \Leftrightarrow \sum_{i,j} \partial_t c_{i,j}(t) \mathbf{u}_{i,j} &= \sum_{i,j} c_{i,j}(t) \lambda_{i,j} \mathbf{u}_{i,j} \end{aligned}$$

Multiplying both sides with $\mathbf{u}_{k,l}$ and exploiting the orthonormality $\langle \mathbf{u}_{i,j}, \mathbf{u}_{k,l} \rangle = \delta_{i,k} \delta_{j,l}$ where δ represents the Kronecker symbol gives us the partial derivative $\partial_t \mathbf{L}$ projected onto eigenvector $\mathbf{u}_{k,l}$. This differential equation can be easily solved for $c(t)$.

$$\begin{aligned} \left\langle \mathbf{u}_{k,l}, \sum_{i,j} \partial_t c_{i,j}(t) \mathbf{u}_{i,j} \right\rangle &= \left\langle \mathbf{u}_{k,l}, \sum_{i,j} c_{i,j}(t) \lambda_{i,j} \mathbf{u}_{i,j} \right\rangle \\ \Leftrightarrow \partial_t c_{k,l}(t) &= c_{k,l}(t) \lambda_{k,l} \\ \Leftrightarrow c_{k,l}(t) &= \exp(\lambda_{k,l} t) c_{k,l}(0) \end{aligned}$$

We finally get an explicit formula for $\mathbf{L}(t)$: The scale space representation \mathbf{L} is the solution of the discretized diffusion equation and has the form

$$\mathbf{L}(t) = \sum_{i,j} c_{i,j}(t) \mathbf{u}_{i,j} = \sum_{i,j} \exp(\lambda_{i,j} t) c_{i,j}(0) \mathbf{u}_{i,j}$$

with scalar coefficients $c_{i,j}(t) = \langle \mathbf{L}(t), \mathbf{u}_{i,j} \rangle$. In matrix representation, the solution simplifies to $\mathbf{L}(t) = \mathbf{U} \exp(\mathbf{\Lambda} t) \mathbf{U}^T \mathbf{L}(0)$. Partial derivatives of any order $\partial_{t^n} \mathbf{L}(t)$ can be easily computed using

$$\partial_{t^n} \mathbf{L}(t) = \sum_{i,j} \partial_{t^n} \exp(\lambda_{i,j} t) c_{i,j}(0) \mathbf{u}_{i,j} = \sum_{i,j} \lambda_{i,j}^n \exp(\lambda_{i,j} t) c_{i,j}(0) \mathbf{u}_{i,j}$$

which, in matrix representation, simplifies to $\mathbf{L}(t) = \mathbf{U} \mathbf{\Lambda}^n \exp(\mathbf{\Lambda} t) \mathbf{U}^T \mathbf{L}(0)$.

5.1 Reduced Eigenimages

Analyzing the time evolution $\mathbf{L}(t) = \sum_{i,j} \exp(\lambda_{i,j} t) c_{i,j}(0) \mathbf{u}_{i,j}$, which is merely a weighted sum of eigenimages, it becomes obvious that eigenimages with smaller eigenvalues have less influence on scales for $t > 0$ than those with eigenvalues near 0.

Omitting less influential eigenimages allows us to reduce memory and time complexity with only moderate changes in the scale space representation. The sum of squared differences SSD between the $\mathbf{L}(t)$ and $\tilde{\mathbf{L}}(t)$ with

$$\tilde{\mathbf{L}}(t) = \sum_{i,j, \lambda_{i,j} > \lambda_{min}} \exp(\lambda_{i,j} t) c_{i,j}(0) \mathbf{u}_{i,j}$$

converges rapidly to 0.




Fig. 2. Lossy compression of signal f . Each image is a reconstruction of f from a reduced set of eigenimages. The number of eigenimages is shown above each image.




Fig. 3. Sampled scale space representation \mathbf{L} (left top), sampled scale space representation $\tilde{\mathbf{L}}$ with reduced eigenimages (left bottom) and the sum of squared differences from low scale to high scale images (right)

This confirms the intuitive expectation that eigenimages with lower eigenvalues correspond to those high frequency components that vanishing very fast with increasing scale. The low frequency components or eigenimages with bigger eigenvalues are more robust against smoothing. Therefore, eliminating eigenimages with small eigenvalues roughly compares to smoothing.

The aforementioned method is comparable to a dimensionality reduction in eigenspace of the Laplacian operator. It is independent of prior information of the content of \mathbf{L} . Taking the initial image $\mathbf{L}(0)$ and thus $c_{i,j}(0)$ into account would allow for further error reduction.

6 Conclusion and Future Work

The discrete scale space as an equivalent to the two-dimensional Gaussian scale space has been discussed and important properties have been derived. A computationally practicable implementation of the discrete scale space framework has been outlined. Our regular 6-neighborhood, a periodic boundary condition and a suitable critical point definition respecting the Euler number [7] have led to discrete scale space axioms that differ from those proposed by Lindeberg. A computationally efficient sampling method, based on properties of the Laplacian kernel. A first investigation of the deep structure of the discrete scale space was illustrated with Laplacian eigenimage in scale space. However, further and more formal investigation of the deep structure (structural changes under the influence of changes) of the discrete scale space is necessary. We note that the decomposition can borrow useful insights from the area of heat kernels on graphs [13]. Since our grid can be considered as a well-structured graph on which we apply a heat kernel-like approach, exploiting such properties [14, 15] under changing scale may give new insights.

References

1. Viola, P., Jones, M.: Robust real-time object detection. *IJCV* (2001)
2. Kuijper, A.: Exploring and exploiting the structure of saddle points in gaussian scale space. *Computer Vision and Image Understanding* 112(3), 337–349 (2008)
3. Lindeberg, T.: *Discrete Scale-Space Theory and the Scale-Space Primal Sketch*. PhD thesis, Royal Institute of Technology (1991)
4. Kuijper, A.: On detecting all saddle points in 2D images. *Pattern Recognition Letters* 25(15), 1665–1672 (2004)
5. Kuijper, A., Florack, L.M.J.: Understanding and Modeling the Evolution of Critical Points under Gaussian Blurring. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002, Part I. LNCS*, vol. 2350, pp. 143–157. Springer, Heidelberg (2002)
6. Kuijper, A., Florack, L.M.J.: The Relevance of Non-generic Events in Scale Space Models. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002, Part I. LNCS*, vol. 2350, pp. 190–204. Springer, Heidelberg (2002)
7. Tschirsich, M., Kuijper, A.: A Discrete Scale Space Neighborhood for Robust Deep Structure Extraction. In: *S+SSPR 2012*, vol. 7626, pp. 124–132. Springer, Heidelberg (2012)
8. Witkin, A.P.: Scale-space filtering. In: *Proc. 8th Int. Joint Conf. Art. Intell.*, pp. 1019–1022. Karlsruhe, Germany (August 1983)
9. Koenderink, J.J.: The structure of images. *Biological Cybernetics* 50, 363–370 (1984)
10. Weickert, J., Ishikawa, S., Imiya, A.: Scale-space has been discovered in Japan. Technical report, Department of Computer Science, University of Copenhagen (August 1997)
11. Iijima, T.: Basic theory on normalization of a pattern (in case of typical one-dimensional pattern). *Bulletin of Electrical Laboratory*, 368–388 (1962)
12. Tschirsich, M.: The discrete scale space as a base for robust scale space algorithms. Technical report, Department of Computer Science, Technical University of Darmstadt (June 2012)
13. Hancock, E.R., Wilson, R.C.: Pattern analysis with graphs: Parallel work at Bern and York. *Pattern Recognition Letters* 33(7), 833–841 (2012)
14. Xiao, B., Hancock, E.R., Wilson, R.C.: Geometric characterization and clustering of graphs using heat kernel embeddings. *Image Vision Comput.* 28(6), 1003–1021 (2010)
15. Xiao, B., Hancock, E.R., Wilson, R.C.: Graph characteristics from the heat kernel trace. *Pattern Recognition* 42(11), 2589–2606 (2009)