Deformable Kernels for Early Vision

Pietro Perona

Abstract—**Early vision algorithms often have a first stage of linear-filtering that 'extracts' from the image information at multiple scales of resolution and multiple orientations. A common difficulty in the design and implementation of such schemes is that one feels compelled to discretize coarsely the space of scales and orientations in order to reduce computation and storage costs. This discretization produces anisotropies due to a loss of translation-, rotation-, and scaling-invariance that makes early vision algorithms less precise and more difficult to design. This need not be so: one can compute and store efficiently the response of families of linear filters defined on a *continuum* of orientations and scales. A technique is presented that allows 1) computing the best approximation of a given family using linear combinations of a small number of 'basis' functions; 2) describing all finite-dimensional families, i.e., the families of filters for which a finite dimensional representation is possible with no error. The technique is based on singular value decomposition and may be applied to generating filters in arbitrary dimensions and subject to arbitrary deformations; the relevant functional analysis results are reviewed and precise conditions for the decomposition to be feasible are stated. Experimental results are presented that demonstrate the applicability of the technique to generating multi-orientation multi-scale 2D edge-detection kernels. The implementation issues are also discussed.**

Index Terms—**Steerable filters, wavelets, early vision, multiresolution image analysis, multirate filtering, deformable filters, scale-space**

I. Introduction

Points, lines, edges, textures, and motions are present in almost all images of the everyday world. These elementary visual structures often encode a great proportion of the information contained in the image, moreover they can be characterized using a small set of parameters that are locally defined: position, orientation, characteristic size or scale, phase, curvature, velocity. It is threrefore reasonable to start visual computations with measurements of these parameters. The earliest stage of visual processing, common for all the classical early vision modules, could consist of a collection of operators that calculate one or more dominant orientations, curvatures, scales, velocities at each point of the image or, alternatively, assign an 'energy,' or 'probability,' value to points of a position-orientation-phase-scale-etc.-space. Ridges and local maxima of this energy would mark special interest loci such as edges and junctions. The idea that biological visual systems might analyze images along dimensions such as orientation and scale dates back to work by Hubel and Wiesel [1], [2] in the 1960s. In the computational vision literature the idea of analyzing images along multiple orientations appears at the beginning of the seventies with the Binford-Horn linefinder [3], [4] and later work by Granlund [5].

A computational framework that may be used to performs this proto-visual analysis is the convolution of the image with kernels of various shapes, orientations, phases, elongation, and scale. This approach is attractive because it is simple to describe, implement, and analyze. It has been proposed and demonstrated for a variety of early vision tasks [6] - [24]. Various 'general' computational justifications have been proposed for basing visual processing on the output of a rich set of linear filters: 1) Koenderink has argued that a structure of this type is an adequate substrate for local geometrical computations [25] on the image brightness, and that it may derived from axioms of invariance with respect to rotation, translation and scaling [26], [27], [28], 2) Adelson and Bergen [24] have derived it from the 'first principle' that the visual system computes derivatives of the image along the dimensions of wavelength, parallax, position, time, 3) a third point of view is the one of 'matched filtering': the kernels are synthesized to match the visual events that one looks for.

The kernels that have been proposed in the computational literature have typically been chosen according to one or more of three classes of criteria: 1) 'generic optimality' (e.g. optimal sampling of space-frequency space [29]), 2) 'task optimality' (e.g., signal to noise ratio, localization of edges [10]), and 3) emulation of biological mechanisms [30], [31], [32]. While there is no general consensus in the literature on precise kernel shapes, there is convergence on kernels roughly shaped like either Gabor functions, or derivatives or differences of either round or elongated Gaussian functions—all these functions have the advantage that they can be specified and computed easily. A good rule of the thumb in the choice of kernels for early vision tasks is that they should have good localization in space and frequency, and should be roughly tuned to the visual events that one wants to analyze.

Since points, edges, lines, textures, and motions can exist at all possible positions, orientations, scales of resolution, and curvatures one would like to be able to use families of filters that are tuned to all orientations, scales, and positions. Therefore, once a particular convolution kernel has been chosen, one would like to convolve the image with deformations (rotations, scalings, stretchings, bendings, etc.) of this 'template.' In reality one can afford only a finite (and small) number of filtering operations, hence the common practice of 'sampling' the set of orientations, scales, positions, curvatures, and phases:

- **Motion flow** computation using spatiotemporal filters has been proposed by Adelson and Bergen [9] as a model of human vision and has been demonstrated by Heeger

Manuscript received Oct. 28, 1991; revised Jan. 24, 1994.

The author is with the Dept. of Engineering and Applied Science 116-81, California Institute of Technology, Pasadena, CA 91125, and Dipartimento di Eletronica ed Informatica, Università di Padova, Italy; e-mail perona@caltech.edu.

IEEECS Log Number P95045.

[11] (his implementation had 12 discrete spatio-temporal orientations and three scales of resolution).

- **Texture** analysis algorithms with multiple-resolution multiple-orientation kernels are due to Knuttson and Granlund [6] (four scales, four orientations, two phases), Turner [12] (four scales, four orientations, two phases), Fogel and Sagi [16] (four scales, four orientations, two phases), Malik and Perona [17] (11 scales, six orientations, one phase) and Bovik et al. [18] (n scales, m orientations, one phase).
- **Stereo** algorithms by Kass [7] (12 filters, scales, orientations, and phases unspecified) and Jones and Malik [33], [34] (six scales, two to six orientations, two phases).
- **Curved line grouping** algorithms by Parent and Zucker [15] (one scale, eight orientations, one phase) and Malik and Gigus [35] (nine curvatures, one scale, 18 orientations, two phases).
- **Brightness boundary detection** work by Binford and Horn [3], [4] (24 orientations), Canny [10] (one to two scales, ∞-six orientations, one phase), Morrone, Owens, and Burr [13], [14] (one to three scales, two to four orientations, ∞ phases), work on edge and illusory contour detection by Heitger, Rosenthaler, Kübler, and von der Heydt [36], [37] (six orientations, one scale, two phases).
- **Image compression** work by Daugman [38] (four scales, six orientations, two phases) and by Zhong and Mallat [19] (four scales, two orientations, one phase).

Discretization has the strong drawback of introducing anisotropies and algorithmic difficulties in the computational implementations. It would be preferable to keep thinking in terms of a continuum, of angles for example, and be able to localize the orientation of an edge with the maximum accuracy allowed by the filter one has chosen.

This aim may sometimes be achieved by means of interpolation: One convolves the image with a small set of kernels, say at a number of discrete orientations, and obtains the result of the convolution at *any* orientation by taking linear combinations of the results. Since convolution is a linear operation the interpolation problem may be formulated in terms of the kernels (for the sake of simplicity the case of rotations in the plane is discussed here): Given a kernel F: $\mathbf{R}^2 \rightarrow \mathbf{C}^2$, define the family of 'rotated' copies of F as: $F_\theta = F \circ R_\theta$, $\theta \in \mathbf{S}^1$, where $\mathbf{S}^1$ is the circle and R_θ is a rotation. Sometimes it is possible to express F_θ as

$$F_\theta(\mathbf{x}) = \sum_{i=1}^{n} \alpha(\theta)_i G_i(\mathbf{x}) \quad \forall \theta \in \mathbf{S}^1, \forall \mathbf{x} \in \mathbf{R}^2 \tag{1}$$

a *finite* linear combination of functions G_i: $\mathbf{R}^2 \rightarrow \mathbf{C}^1$. It must be noted that, at least for positions and phases, the mechanism for realizing this in a systematic way is well understood: in the case of positions the sampling theorem gives conditions and an interpolation technique for calculating the value of the filtered image at any point in a continuum; in the case of phases a pair of filters in quadrature can be used for calculating the response at any phase [9], [39]. Rotation, scalings, and other deformations are less well understood (schemes for estimating prevalent orientation in a continuum have been proposed [40], [41], [42]).

An example of 'rotating' families of kernels that have a finite representation is well known: the first derivative along an arbitrary direction of a round ($\sigma_x = \sigma_y$) Gaussian may be obtained by linear combination of the X- and Y-derivatives of the same (this property is of course common to all rotationally-simmetric functions). Most implementations of the Canny edge detector [10] are based on this principle. Unfortunately, the kernel obtained this way has poor orientation selectivity and therefore it is not suitable for edge detection if one wants to recover edge-junctions. Freeman and Adelson have first proposed [43] that it would be desirable to construct orientation-selective kernels that can be exactly rotated by interpolation (they call this property "steerability" and the term will be used in this paper) and have shown that higher order derivatives of round Gaussians, indeed all polynomials multiplied by a radially symmetric function are steerable (they have a more general class of functions [44]—it is described in the comments to Theorem 1). For high polynomial orders these functions may be designed to have higher orientation selectivity and can be used for contour detection and signal processing [22], [45]. This line of work has been extended to scale by Simoncelli et al. in [48]. However, one must be aware of the fact that for most kernels F of interest a finite decomposition of F_θ as in (1) cannot be found. For example the elongated kernels used in edge detection by [46], [23], [47] do not have a finite decomposition as in (1).

Lenz [27] has independently proposed using group-theoretical methods for steering a given filter; his analysis and technique are optimal and applicable whenever the underlying deformations have group structure (i.e., rotations). As will be discussed later, some deformations (e.g., scale) do not have group structure and, therefore, his technique is not always applicable.

One needs an approximation technique that, given an F_θ, allows one to generate a function $G_\theta^{[n]}$ which is sufficiently similar to F_θ and that is steerable, i.e., can be expressed as a finite sum of n terms as in (1). In the case of rotations, Freeman and Adelson propose to approximate the kernel with an adequately high order polynomial multiplied by a radially symmetric function (which they show is steerable). However, this method does not guarantee a parsimonious approximation: given a tolerable amount of error one would like to find an approximating $G_\theta^{[n]}$ that has minimum number n of components. A different design perspective could also be taken: given a number n of filtering operations allowed, synthesize the best (with respect to the specific task at hand) kernel *within* the class of functions that can be *exactly* represented by a sum of n terms. Therefore it is useful to be able to answer to the question: What is the set of functions that can be represented exactly as in (1)? This paper addresses these questions in detail. Partial results have been reported in [49], [50], [51], [52].

This paper is organized as follows: the special case of the rotations (1) is explored and solved in Section II and the appendix. In Section III a few results from functional analysis are

recalled to extend the approximation technique to all 'compact' deformations. In Section IV an application of the approximation technique to generating steerable filters for edge detection is described. In Section V it is shown how to generate a steerable and scalable family. Experimental results and implementation issues are presented and discussed for the schema presented in Sections V and VI.

II. Steerable Approximations

In order to solve the approximation problem proposed in the introduction one needs of course to define the 'quality' of the approximation $G_\theta^{[n]} \approx F_\theta$. There are two reasonable choices: 1) a distance $D(F_\theta, G_\theta^{[n]})$ in the space $\mathbf{R}^2 \times \mathbf{S}^1$ where F_θ is defined; 2) if F_θ is the kernel of some filter one is interested in the worst-case error in the 'output' space: the maximum distance $d(\langle F_\theta, f\rangle, \langle G_\theta^{[n]}, f\rangle)$ over all unit-norm f defined on $\mathbf{R}^2$ (the inner products are, of course, taken in $L^2(\mathbf{R}^2)$ and the distance is on the Banach space of functions of $\mathbf{S}^1$). The symbols Δ_n and δ_n will indicate the 'optimal' distances, i.e., the minimum possible approximation errors using n components. Among many reasonable definitions of distance, the one induced by the L^2-norm has the advantage that it makes calculations particularly convenient. This is the definition that will be used in this study:

Definition.

$$D_n\left(F_\theta, G_\theta^{[n]}\right) = \left\|F_\theta - G_\theta^{[n]}\right\|_L^2\left(\mathbf{R}^2 \times \mathbf{S}^1\right)$$

$$\Delta_n(F_\theta) = \inf_{G_\theta^{[n]}} D_n\left(F_\theta, G_\theta^{[n]}\right)$$

$$d_n\left(F_\theta, G_\theta^{[n]}\right) = \sup_{\|f\|=1}\left\|\left\langle F_\theta - G_\theta^{[n]}, f\right\rangle_{\mathbf{R}^2}\right\|_L^2\left(\mathbf{S}^1\right)$$

$$\delta_n(F_\theta) = \inf_{G_\theta^{[n]}} d_n\left(F_\theta, G_\theta^{[n]}\right)$$

Consider now the approximation to F_θ defined as follows:

Definition. Call $F_\theta^{[n]}$ the n-terms sum:

$$F_\theta^{[n]} = \sum_{i=1}^{n} \sigma_i a_i(\mathbf{x}) b_i(\theta) \tag{2}$$

with σ_i, a_i and b_i defined in the following way: let $\hat{h}(\nu)$ be the Fourier series expansion of the function $h(\theta)$ defined by:

$$h(\theta) = \int_{\mathbf{R}^2} F_\theta(\mathbf{x})\overline{F_{\theta'=0}(\mathbf{x})}\,d\mathbf{x} \tag{3}$$

(any fixed value is fine for θ', notice that $F_{\theta'=0}(\mathbf{x}) = F(\mathbf{x})$) and let ν_i be the frequencies on which $\hat{h}(\nu)$ is defined, ordered in such a way that $\hat{h}(\nu_i) \geq \hat{h}(\nu_j)$ if $i \leq j$. Call $N \leq \infty$ the number of nonzero terms $\hat{h}(\nu_i)$—it turns out that $\hat{h}(\nu_i) \geq 0$ (see Theorem 1 and proof in the appendix). Finally for $i \leq N$ define the quantities:

$$\sigma_i = \hat{h}(\nu_i)^{1/2} \tag{4}$$

$$b_i(\theta) = e^{j2\pi\nu_i\theta} \tag{5}$$

$$a_i(\mathbf{x}) = \sigma_i^{-1}\int_{\mathbf{S}^1} \overline{F_\theta(\mathbf{x})}e^{j2\pi\nu_i\theta}\,d\theta \tag{6}$$

Then $F_\theta^{[n]}$ is the best n-dimensional approximation to F_θ in the following sense:

Theorem 1. *Given the definitions and notation introduced above, suppose that $F \in L^2(\mathbf{R}^2)$ then:*

1) *$\{a_i\}$ and $\{b_i\}$ are orthonormal sequences of functions in $L^2(\mathbf{R}^2)$ and $L^2(\mathbf{S}^1)$, respectively.*
2) *$F_\theta^{[n]}$ is the smallest possible exact representation of F_θ, i.e., if $\exists M$, β_i, g_i s.t. $F_\theta(\mathbf{x}) = \sum_{i=1}^{M}\beta_i(\theta)g_i(\mathbf{x})$ then $M \geq N$.*
3) *The number N of terms is finite iff the number M of indices i for which $a_i(\mathbf{x}) \neq 0_{L^2(\mathbf{R}^2)}$ is finite, and $N = M$.*
4) *$F_\theta^{[n]}$ is an optimal n-approximation of F_θ with respect to both distances:*

$$D_n\left(F_\theta, F_\theta^{[n]}\right) = \Delta_n(F_\theta) = \left(\sum_{i=n+1}^{N}\sigma_i^2\right)^{1/2} \tag{7}$$

$$d_n\left(F_\theta, F_\theta^{[n]}\right) = \delta_n(F_\theta) = \sigma_{n+1} \tag{8}$$

5) *D_n, $\delta_n \to 0$ for $n \to N$.*
6) *$F_\theta^{[n]} = F_0^{[n]} \circ R_\theta$.*
7) *$\forall\theta_1, \ldots, \theta_n$ with the exception of a set of measure zero $\exists b_1', \ldots, b_n'$ s.t. $F_\theta^{[n]}(x) = \sum_{i=1}^{n} F_{\theta_i}^{[n]}(\xi)b_n'(\theta)$.*

Comment.

1. The expression for the b_i is independent of F. Only σ_i and a_i depend on F. The b_i depend on the particular group of transformations (the rotations of the plane in this case) that is used to generate F_θ from F.
3. The 'if' part of statement 3 is equivalent to the main theorem of [44]—Freeman and Adelson show that all functions that one can write as a linear combination of functions like the a_is (polar-separable with sinusoidal θ component) are steerable. The 'only if' part says that the functions that they described are *all* the steerable functions.
4. For deciding at what point n to truncate the sum one plots the error δ_n or Δ_n vs. n and looks for the smallest integer n for which the error is less than some assigned value. See Figs. 2 and 3.
6. This means that $F_\theta^{[n]}$is steerable, i.e., its shape does not change with θ, modulo a rotation in the domain. Therefore $F_\theta^{[n]}$is the best approximation to F_θ in the space of 'n-steerable' functions ('best' is intended with respect to the L^2-induced distance).

7.1. A set of size n of rotated copies of $F_\theta^{[n]}$is enough for representing $F_\theta^{[n]}$, so we may choose to use this different de-

composition instead of 2. On the other hand, this representation has some numerical disadvantages:

1) The set F_{θ_i} is not orthonormal, so its numerical implementation is less efficient (it will require more significant bits in the calculations to obtain the same final precision).
2) The functions a_i are easier to approximate with sums of X-Y separable functions than the F_{θ_i} (see the experimental Section IV.B, and Fig. 5).

7.2. The error $d_n(F_\theta, F_\theta^{[n]})$ of the n-approximation is constant with respect to θ since $F_\theta = F \circ R_\theta$ and $F_\theta^{[n]} = F_0^{[n]} \circ R_\theta$. There is *no anisotropy* even if $F^{[n]}$ is an approximation.

The proof of this theorem is reported in the appendix. It is based on the fact that the triples $(\sigma_i, a_i(\mathbf{x}), b_i(\theta))$ are the singular value decomposition (SVD) of a linear operator associated to $F_\theta(\mathbf{x})$. (From now on the abbreviation SVD will be used to indicate the decomposition of a kernel into such triples).

III. Deformable Functions

The existence of the optimal finite-sum approximation of the kernel $F_\theta(\mathbf{x})$ as described in the previous section and Appendix A is not peculiar to the case of rotations. This is true in more general circumstances: this section collects a few facts of functional analysis that show that one can compute finite optimal approximations to continuous families of kernels whenever certain 'compactness' conditions are met.

Consider a parametrized family of kernels $F(\mathbf{x}; \theta)$ where $\mathbf{x} \in X$ now indicates a generic vector of variables in a set X and $\theta \in T$ a vector of parameters in a set T. (The notation is changed slightly from the previous section.) Consider the sets A and B of continuous functions from X and T to the complex numbers, call $a(\mathbf{x})$ and $b(\theta)$ the generic elements of these two sets. Proceed as at the beginning of the appendix and consider the operator $L : A \rightarrow B$ defined by F as $(La(\cdot))(\theta) = \langle F(\cdot\,; \theta), a(\cdot)\rangle_A$.

A first theorem says that if the kernel F has bounded norm then the associated operator L is compact (see [53]):

THEOREM 2. *Let X and T be locally compact Hausdorff spaces and $F \in L_2(X \times T)$. Then L is well defined and is a compact operator.*

Such a kernel is commonly called a *Hilbert-Schmidt kernel.*

A second result tells us that if a linear operator is compact, then it has a discrete spectrum (see [54]):

THEOREM 3. *Let L be a compact operator on (complex) normed spaces, then the spectrum S of L is at most denumerable.*

A third result says that if L is continuous and operates on Hilbert spaces then the compactness property transfers to the adjoint of L (see [54]):

THEOREM 4. *Let L be a compact operator on Hilbert spaces, then the adjoint L^* is compact.*

The composition of two compact operators is compact, so the operators LL^* and L^*L are compact and will have a discrete spectrum as guaranteed by Theorem 3. The singular value decomposition (SVD) for the operator L can therefore be computed as the collection of triples (σ_i, a_i, b_i), $i = 0, \ldots$ where the σ_i constitute the spectra of both LL^* and L^*L and the a_i and b_i are the corresponding eigenvectors.

The last result can now be enunciated (see [55]):

THEOREM 5. *Let $L : A \rightarrow B$ be a linear compact operator between two Hilbert spaces. Let a_i, b_i, σ_i be the singular value decomposition of L, where the σ_i are in decreasing order of magnitude. Then*

1) *An optimal n-dimensional approximation to L is* $L_n = \sum_{i=1}^{n} \sigma_i a_i b_i$

2) *The approximation error is* $\delta_n(L) = \sigma_{n+1}$, $\Delta_n^2(L) = \sum_{i=n+1}^{N} \sigma_i^2$.

As a result we know that when our original template kernel $F(\mathbf{x})$ and the chosen family of deformations $R(\theta)$ define a Hilbert-Schmidt kernel $F(\mathbf{x}; \theta) = (F \circ R(\theta))(\mathbf{x})$ then it is possible to compute a finite discrete approximation as for the case of 2D rotations.

Are the families of kernels $F(\mathbf{x}; \theta)$ of interest in vision Hilbert-Schmidt kernels? In the cases of interest for vision applications the 'template' kernel $F(\mathbf{x})$ typically has a finite norm, i.e., it belongs to $L_2(X)$ (all kernels used in vision are bounded compact-support kernels such as Gaussian derivatives, Gabors, etc.). However, this is not a sufficient condition for the family $F(\mathbf{x}; \theta) = F \circ R(\theta)$ $(\mathbf{x})$ obtained composing $F(\mathbf{x})$ with deformations $R(\theta)$ (rotations, scalings) to be a Hilbert-Schmidt kernel: the norm of $F(\mathbf{x}; \theta)$ could be unbounded. A sufficient condition for the associated family $F(\mathbf{x}; \theta)$ to be a Hilbert-Schmidt kernel is that the inverse of the Jacobian of the transformation R, $|JR|^{-1}$ belongs to $L_2(T)$. In fact the norm of $F(\cdot\,; \cdot)$ is bounded above by the product of the norm of F in X and the norm of $|JR|^{-1}$ in T:

$$\begin{aligned}\|F(\cdot;\cdot)\|^2 &= \int_T \int_X |F(\mathbf{x}; \theta)|^2 \, d\mathbf{x} d\theta \\ &= \int_T \int_X |(F \circ R(\theta))(\mathbf{x})|^2 \, d\mathbf{x} d\theta \\ &\leq \int_T \int_X |JR(\theta)|^{-1} |F(\mathbf{y})|^2 \, d\mathbf{y} d\theta \\ &= \|F(\cdot)\|^2 \left\| |JR(\cdot)|^{-\frac{1}{2}} \right\|^2\end{aligned}$$

which is bounded by hypothesis.

A typical condition in which this arises is when the transformation R is unitary, e.g., a rotation, a translation, or an appropriately normalized scaling, and the set T is bounded. In that case the norm of $\|JR\|^{-1}$ is equal to the measure of T. The following sections in this paper will illustrate the power of these results by applying them to the decomposition of rotating 2D kernels (Section II), 2D kernels into sums of X-Y-separable kernels (Section IV.B), and rotating and scaled kernels (Section V).

A useful subclass of kernels F for which the finite orthonormal approximation can be in part explicitly computed is obtained by composing a template function with transforma-

tions T_θ belonging to a compact group. This situation arises in the case of n-dimensional rotations and is useful for edge detection in tomographic data and spatiotemporal filtering [27], [49], [50] (the book by Lenz [27] contains a very useful introduction to the group-theoretical concepts that are needed).

IV. Steerable Approximations: Practical Issues and Experimental Results

In this section the formalism described so far is applied to the problem of generating steerable and scalable approximations to convolution kernels for an edge-detector. The Gaussian-derivative kernels used by [46], [23] have been chosen for this example. These kernels have an elongated shape to achieve high orientation selectivity (equivalently, narrow orientation tuning). The real part of the kernels is a Gaussian $G(\mathbf{x}, \sigma_x, \sigma_y) = \exp - ((x/\sigma_x)^2 + (y/\sigma_y)^2)$ differentiated twice along the Y axis. The imaginary part is the Hilbert transform of the real part taken along the Y axis (see Figs. 1 and 3).

One more twist is added: the functions a_i in the decomposition may in turn be decomposed as sums of a small number of X-Y-separable functions making the implementation of the filters considerably faster.

A. Rotations

In the case of rotations Theorem 1 may be used directly to compute the decomposition. The calculations proceded as indicated in Section II. For convenience they are summarized in a recipe:

1) Select the 'template' kernel $F(\mathbf{x})$ of which one wants rotated versions $F_\theta(\mathbf{x}) = F \circ R_\theta(\mathbf{x}) = F(x\cos(\theta) + y\sin(\theta), -x\sin(\theta) + y\cos(\theta))$ and the maximum tolerable error η.
2) Compute numerically the function $h(\theta)$ using its definition (3). The sampling period should be a fraction of the orientation selectivity of the kernel.
3) Compute the discrete Fourier transform $\hat{h}$ of h. Verify that the coefficients $\hat{h}_k$ are non-negative.
4) Order the $\hat{h}_k$ computed at the previous step by decreasing magnitude and call their square roots σ_i (see Fig. 1) and the corresponding frequencies ν_i (see (4)).
5) Define the functions $b_i(\theta)$ according to (5) and the ν_i calculated at the previous step.
6) Compute the error plots $\delta(n)$ and $\Delta(n)$ from (7) and (8) (see Fig. 3). Obtain the number n of components required for the approximation as the first integer where the error drops below the tolerable error η.
7) Compute the functions $a_i(\mathbf{x})$ using (6). (See Fig. 1).
8) The n-approximation of $F_\theta(\mathbf{x})$ can now be calculated using (2).

The numerical implementation of the formulae of Section II is straightforward. In the implementation used to produce the figures and the data reported in this section the kernels F_θ were defined on a 128×128 array of single-precision floating-point numbers. The set of all angles was discretized in 128 samples. The Y-axis variance was $\sigma_y = 8$ pixels, and the X-axis variance was $\sigma_x = k\sigma_y$ with $k = 1, 2, 3$. Calculations were reduced by a factor of two exploiting the hermitian symmetry of these kernels; the number of components can also be halved—the experimental data given below and in the figures are calculated this way.

Notice first that the coefficients σ_i converge to zero exponentially fast; as a consequence the same is true for both errors. This is very important in practice since it implies that a very small number of coefficients is required. In Fig. 2 the approximate reconstruction is shown for $n = 4, 9, 15$. Notice that the elongation and therefore the 'orientation selectivity' of the filter increases with the number of components. In Fig. 4 the modulus of the response of the complex filter to an edge is shown for two different kernels and increasing levels of approximation. The number n of singular components required to reconstruct the $\sigma_x : \sigma_y = 1{:}1$, and $\sigma_x : \sigma_y = 1{:}2$ families is smaller as indicated by the plots and in the caption of Fig. 3.

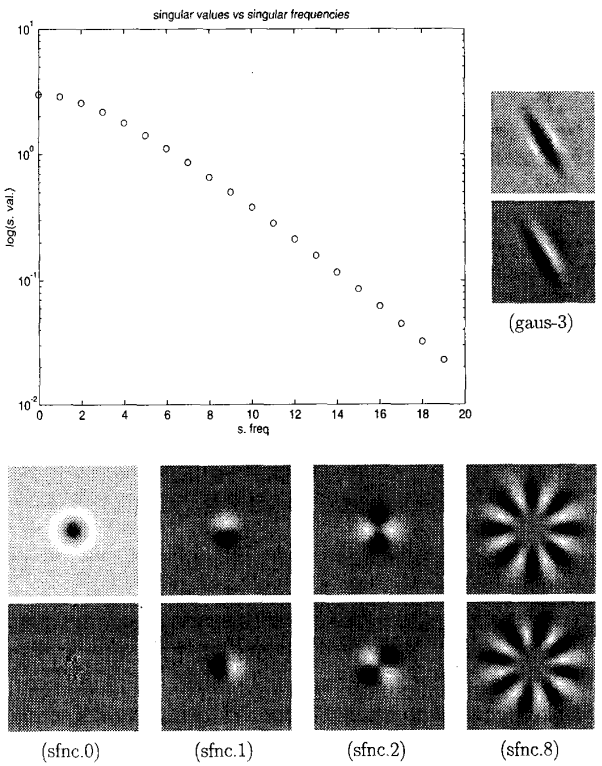

Fig. 1. The decomposition (a_i, b_i, σ_i) of a complex kernel used for brightness-edge detection [23]. (Top right) The template function (gaus-3) is shown rotated counterclockwise by 120°. Its real part (above) is the second derivative along the vertical (Y) axis of a Gaussian with $\sigma_x : \sigma_y$ ratio of 1 : 3. The imaginary part (below) is the Hilbert transform of the real part along the Y axis. (Top left) The first 20 σ_i (s. val) shown on a logarithmic scale plotted against the associated frequencies $\nu_i = i$ (s. freq). They decay roughly exponentially: $\sigma_{i+1} \approx 0.75\ \sigma_i$. (Bottom) The functions a_i (sfnc.i) are shown for $i = 0, 1, 2, 8$. The real part is above; the imaginary part below. The functions $b_i(\theta)$ are complex exponentials (see text) with associated frequencies $\nu_i = i$.

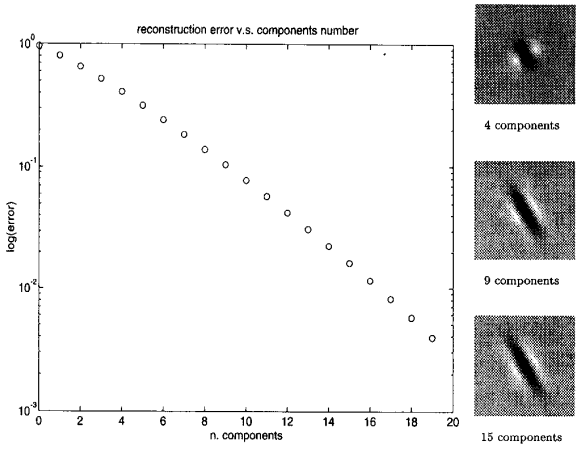

Fig. 2. Approximate reconstruction of the (gaus-3) kernel of Fig. 1. (Right) Reconstruction of the kernel with four components, with nine components, and with 15 components (only real part shown). The optimal reconstruction error Δ_n for $n = 4, 9, 15$ calculated from the singular values (plot on left) is, 52%, 13.9%, and 2.2%, respectively. The reconstruction error $\|\text{gaus-3} - \text{rec.i}\| \cdot \|\text{gaus-3}\|^{-1}$ as measured on the reconstruction in the computer implementation is 50.2%, 13.4%, and 2.2%, respectively (the reconstruction error is independent of the angle and the angle may be chosen in a continuum). The error decreases exponentially, approximately 25-30%, for each component added. (See also the caption of Fig. 3.)

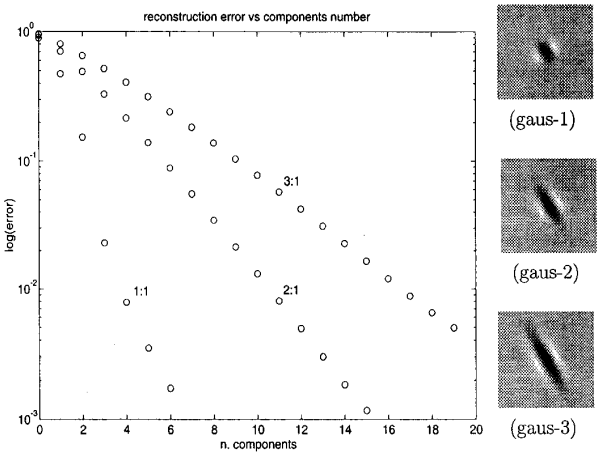

Fig. 3. Comparison of the error plots for three kernels constructed as explained in Fig. 1 from three Gaussians of different aspect ratios. (Right) The real parts of the three kernels shown at an angle of 120°; the ratios $\sigma_x : \sigma_y$ are 1 : 1, 1 : 2, 1 : 3, respectively. (Left) The log of the reconstruction errors are plotted against the number of components employed. For 10% reconstruction error 3, 6, 10 components are needed. For 5% reconstruction error 3, 7, 12 components are needed. Notice that for these Gaussian-derivative functions the reconstruction error decreases roughly exponentially with respect to the number n of components employed: $\Delta_n \approx \exp(-\frac{n}{\tau})$ with $\tau \approx 1.7, 5.2, 8.2$.

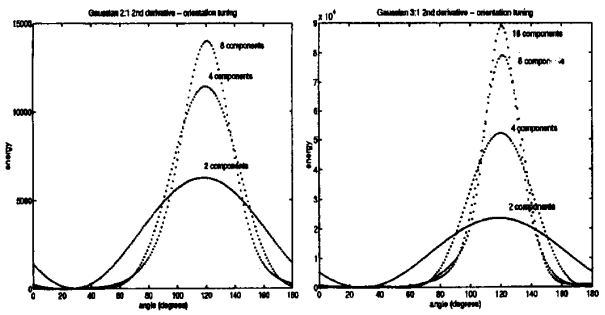

Fig. 4. Magnitude of the response vs. orientation of orientation-selective filters to the image of an edge oriented at 120° (notice the corresponding peak in the filter responses). The kernels of the filters are as in Fig. 3, derived from a 2:1 Gaussian (left) and a 3:1 Gaussian (right). The plots show the response of the filters for different approximations. The first apporximation (two components) gives a broadly tuned response, while the other approximations (4, 8, 12 components) have more or less the same orientation selectivity (half-width of the peak at half-height). The peak of the response sharpens and the precision of the approximation is increased (to 1-2% error for the top curves) when more components are used.

B. X-Y Separability

Whenever a function F is to be used as a kernel for a 2D convolution it is of practical importance to know wether the function is X-Y-separable, i.e., if there are two functions f^x and f^y such that $F(x, y) = f^x(x)f^y(y)$. If this is true the 2D convolution can be implemented cheaply as a sequence of two 1D convolutions.

Even when the kernel F is not X-Y-separable it may be the sum of a small number of separable kernels [56], [57], [58]: $F(x, y) = \sum_i f_i^x(x) f_i^y(y)$. One may notice the analogy of this decomposition with the one expressed in (1): the function $F(x, y)$ may be thought of as a kernel defining an operator from a space A of functions $f^x(x)$ to a space B of functions $f^y(y)$. At a glance one can see that the kernels a_i of Fig. 1 are Hilbert-Schmidt: they have bounded norm. Therefore (see Section III) a discrete optimal decomposition and approximation are possible. Again the SVD machine may be applied to calculate the decomposition of each one of the a_i and its optimal n_i-component approximation. If the SVD of a_i is indicated as: $a_i(x, y) = \sum_{h=1}^{n_i} \rho_{ih} a_{ih}^x(x) a_{ih}^y(y)$ then the approximate decomposition of $F_\theta(x, y)$ expressed in (2) and (20) becomes:

$$F_\theta(x, y) = \sum_{i=1}^{N} \sigma_i b_i(\theta) \sum_{h=1}^{n_i} \rho_{ih} a_{ih}^x(x) a_{ih}^y(y). \qquad (9)$$

How is this done in practice? For all practical purposes the kernels a_i are defined on a discrete lattice. The SVD of a kernel defined on a discrete $M \times N$ rectangular lattice may be computed numerically using any one of the common numerical libraries [59], [60] as if it were an $M \times N$ square matrix A of rank R. The typical notation for the matrix SVD is: $A = UWV^T$ where U and V are orthonormal $M \times R$ and $R \times N$ matrices and

W is an $R \times R$ band-diagonal with positive entries w_i of decreasing magnitude along the diagonal. If this is written as

$$A = \sum_{k=1}^{R} w_k U_k V_k^T \tag{10}$$

where U_k and V_k are columns of U and V the analogy with (2) becomes obvious. Notice that the (hidden in the vector notation) row index of U plays the role of the coordinate y and the row index of V plays the role of the coordinate x. Rewriting the above in continuous notation we obtain:

$$a(x, y) = \sum_{k=1}^{R} w_k u_k(y) v_k(x) \tag{11}$$

The first two terms of the separable decompositions are shown in Fig. 5 for the functions a_3 and a_8.

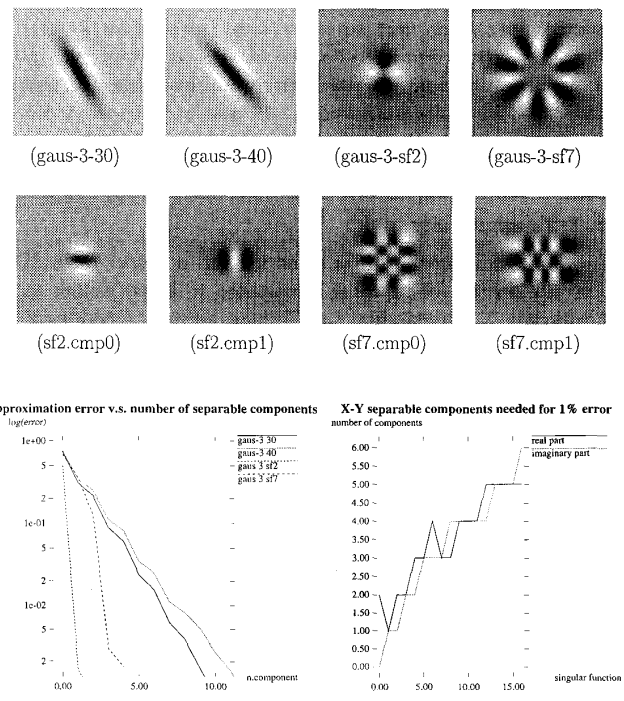

Figure 5:

Fig. 5. Comparison of the 'separability' of two rotated copies of the template function and two elements of the decomposition. The template function is (gaus-3) as in Fig. 1. (Top) The rotation angles are 120° (gaus-3-30) and 130° (gaus-3-40). The functions from the decomposition are a_3 (gaus-3-sf2) and a_8 (gaus-3-sf7) (cfr. Fig. 1). The first two separable components of a_3 are (sf2.cmp0) (sf2.cmp1) and of a_8 are (sf7.cmp0) and (sf7.cmp1). (Bottom left) Approximation error vs. number of separable components compared for the four functions. The rotated copies of the template kernel require more tems than the functions a_i. The number of X-Y separable components necessary to achieve 1% error is: 7 and 8 for the rotated copies of the template function, and 1 and 3 for the singular functions. (Bottom right) The number of components necessary to approximate a_i to less than 1% error is plotted against i. From the plots one may deduce that this number is approximately equal to $1 + i/4$, so that the total number of 1D convolutions required to implement the n-approximation is $2n + n^2/4$.

Whether few or many components will be needed for obtaining a good approximation is again an empirical issue and will depend on the kernel in question. The decomposition of the singular functions a_i associated to the Gaussian-derivative functions used for these simulations is particularly advantageous; the approximation error typically shows a steep drop after a few components are added. This can be seen from the curves in Fig. 5 (bottom-left) where the log of the error is plotted against the number of X-Y-separable components. All the a_i of Fig. 1 can be decomposed this way in sums of X-Y-separable kernels. The number of components needed for approximating each with 1% accuracy or better is indicated in the plots of Fig. 5 (bottom-right) the real and imaginary parts have roughly the same separability. One can see that the number of components increases linearly with i. The caption of Fig. 5 gives precise figures for the Gaussian 3 : 1 case.

It is important to notice that rotated versions of the original template functions F cannot be represented by sums of X-Y-separable functions with the same parsimony (see again Fig. 5 (bottom-left) upper curves). This is one more reason to represent $F_\theta^{[n]}$ as a sum of orthonormal singular functions, rather than as as sum of rotated copies of the template function (Theorem 1, statement 7.), as discussed at the end of Section II. One must remember that beyond X-Y-separation there are a number of techniques for speeding up 2D FIR filtering for example small generating kernel (SGK) filtering [61], that could further speed up the convolutions necessary to implement deformable filtering.

V. Rotation and Scale

A number of filter-based early vision and signal processing algorithms analyze the image at multiple scales of resolution. Although most of the algorithms are defined on, and would take advantage of, the availability of a continuum of scales only a discrete and small set of scales is usually employed due to the computational costs involved with filtering and storing images. The problem of multi-scale filtering is somewhat analogous to the multi-orientation filtering problem that has been analyzed so far: given a template function $F(\mathbf{x})$ and defined $F_\sigma(\mathbf{x})$ as $F_\sigma(\mathbf{x}) = \sigma^{1/2} F(\sigma\mathbf{x})$, $\sigma \in (0, \infty)$ one would like to be able to write F_σ as a (small) linear combination:

$$F_\sigma(\mathbf{x}) = \sum_i s_i(\sigma) d_i(\mathbf{x}) \qquad \sigma \in (0, \infty) \tag{12}$$

Unfortunately the domain of definition of s is not bounded (it is the real line) and therefore the kernel $F_\sigma(\mathbf{x})$ is not Hilbert-Schmidt (it has infinite norm). As a consequence the spectrum of the LL^* and L^*L operators is continuous and no discrete approximation may be computed.

One has therefore to renounce to the idea of generating a continuum of scales spanning the whole positive line. This is not a great loss: the range of scales of interest is never the entire real line. An interval of scales (σ_1, σ_2), with $0 < \sigma_1 \leq \sigma_2 < \infty$ is a very realistic scenario; if one takes the human visual system as an example, the range of frequencies to which it is most sensitive goes from approximately two to 16 cycles per

degree of visual angle i.e., a range of three octaves. In this case the interval of scales is compact and therefore $F_\sigma(\mathbf{x})$ restricted to this interval is Hilbert-Schmidt and one can apply the results of Section III and calculate the SVD and therefore an L_2-optimal finite approximation.

In this section an optimal scheme for doing so is proposed. The problem of simultaneously steering and scaling a given kernel $F(\mathbf{x})$ generating a family $F_{(\sigma,\theta)}(\mathbf{x})$ which has a finite approximation will be tackled. Previous non-optimal schemes are due to Perona [49], [50] and Freeman and Adelson [22].

A. Polar-Separable Decomposition

Observe first that the functions a_i defined in (6) are polar-separable. In fact $\mathbf{x}$ may be written in polar coordinates as $\mathbf{x} = \|\mathbf{x}\| R_{\phi(\mathbf{x})}\mathbf{u}$ where $\mathbf{u}$ is some fixed unit vector (e.g. the first coordinate axis versor) and $\phi(\mathbf{x})$ is the angle between $\mathbf{x}$ and $\mathbf{u}$ and $R_{\phi(\mathbf{x})}$ is a rotation by ϕ. Substituting the definition of F_θ in (6) we get:

$$a_i(\mathbf{x}) = \sigma_i^{-1} \int_{S^1} \overline{F\left(\|\mathbf{x}\| R_{\theta+\phi(\mathbf{x})}(\mathbf{u})\right)} e^{j2\pi\nu_i\theta} d\theta$$

$$= \sigma_i^{-1} e^{-j2\pi\nu_i\phi(\mathbf{x})} \int_{S^1} \overline{F\left(\|\mathbf{x}\| R_{\psi}(\mathbf{u})\right)} e^{j2\pi\nu_i\psi} d\psi$$

so that (2) may be also written as :

$$F_\theta(\mathbf{x}) = \sum_{i=1}^{N} \sigma_i^{-1} c_i(\|x\|) e^{j2\pi\nu_i(\theta-\phi(\mathbf{x}))} \tag{13}$$

$$c_i(\|\mathbf{x}\|) = \sigma_i \int_{S^1} \overline{F\left(\|\mathbf{x}\| R_{\psi(\mathbf{u})}\right)} e^{j2\pi\nu_i\psi} d\psi \tag{14}$$

The scaling operation only affects the radial components c_i and does not affect the angular components. The problem of scaling the kernels a_i, and therefore F_θ through its decomposition, is then the problem of finding a finite (approximate) decomposition of continuously scaled versions of functions $c(\rho)$:

$$c_\sigma(\rho) = \sum_k s_k(\sigma) r_k(\rho) \qquad \sigma \in (\sigma_1, \sigma_2) \tag{15}$$

If the scale interval (σ_1, σ_2) and the function c are such that the operator L associated to F is compact then we can obtain the optimal finite decomposition via the singular value decomposition. The conditions for compactness of L are easily met in the cases of practical importance: it is sufficient that the interval (σ_1, σ_2) is bounded and that the norm of $c(\rho)$ is bounded ($\rho \in \mathbf{R}^+$).

Even if these conditions are met, the calculations usually cannot be performed analytically. One can employ a numerical routine as in Section IV.B for X-Y-separation and for each c_i (below indicated as c^i) obtain an SVD expansion of the form:

$$c^i_\sigma(\rho) = \sum_k \gamma^i_k s^i_k(\sigma) r^i_k(\rho) \tag{16}$$

As discussed before one can calculate the approximation error from the sequence of the singular values γ^i_k. Finally, substituting (16) into (13) and (14) the scale-orientation expansion takes the form (see Fig. 8):

$$F_{\theta,\sigma}(\mathbf{x}) = \sum_{i=1}^{N} \sigma_i e^{j2\pi\nu_i(\theta-\phi(\mathbf{x}))} \sum_{k=1}^{n_i} \gamma^i_k s^i_k(\sigma) r^i_k(\|\mathbf{x}\|) \tag{17}$$

Filtering an image I with a deformable kernel built this way proceeds as follows: First the image is filtered with kernels $a^i_k(\mathbf{x}) = \exp(-j2\pi\nu_i\phi(\mathbf{x})) r^i_k(\|\mathbf{x}\|)$, $i = 0, \ldots, N$, $k = 0, \ldots, n_i$, the outputs I^i_k of this operation can be combined as

$$I_{\theta,\sigma}(\mathbf{x}) = \sum_{i=1}^{N} \sigma_i b_i(\theta) \sum_{k=1}^{n_i} \gamma^i_k s^i_k(\sigma) I^i_k(\mathbf{x})$$

to yield the result. The filtering operations described above can of course be implemented as X-Y-separable convolutions as described in Section IV.B.

A.1. Polar-Separable Decomposition, Experimental Results

An orientation-scale decomposition was performed on the usual kernel (second derivative of a Gaussian and its Hilbert transform, $\sigma_x : \sigma_y = 3 : 1$). The decomposition described in Section IV.A was taken as a starting point. The corresponding functions c_i of (13) are shown in Fig. 6.

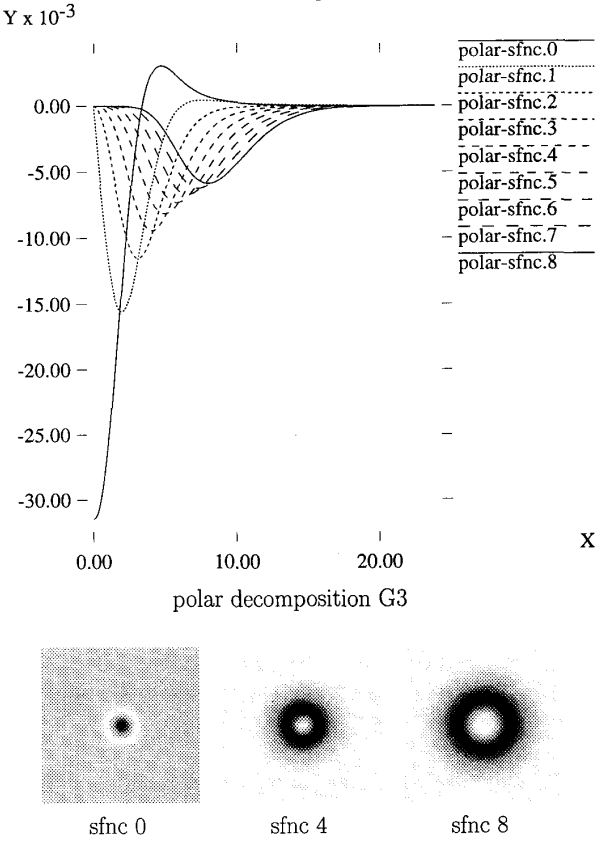

Fig. 6. (Top) The plots of $c_i(\rho)$, the radical part of the singular functions a_i (cfr. (13)). The θ part is always a complex exponential. The original kernel is the same as in Fig. 1. (Bottom) The *0th*, *4th*, and *8th* components c_0, c_4, and c_8 represented in two dimensions.

The interval of scales chosen was (σ_1, σ_2) s.t. $\sigma_1 : \sigma_2 = 1 : 8$, an interval which is arguably ample enough for a wide range of visual tasks.

The range of scales was discretized in 128 samples for computing numerically the singular value decomposition $(\gamma_k^i, s_k^i, r_k^i)$ of $c_\sigma^i(\rho)$. The computed weights γ_k^i are plotted on a logarithmic scale in Fig. 7 (top left). The 'X' axis corresponds to the k index, each curve is indexed by i, $i = 0, \ldots, 8$. One can see that for all the c_i the error decreases exponentially at approximately the same rate. The components $r_k^i(\rho)$ and $s_k^i(\sigma)$, $i = 4$, $k = 0, \ldots, 3$ are shown in the two plots at the bottom of Fig. 7.

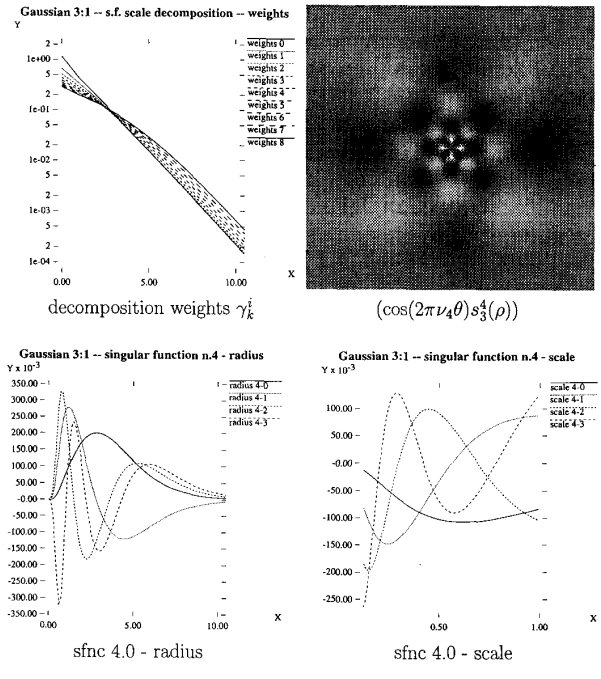

Fig. 7. Scale-decomposition of the radial component of the functions a_i. The interval of scales σ is $\sigma \in (0.125, 1.00)$. See also Fig. 8. (Top left) The weights γ_k^i of each polar function's decomposition ($i = 0, \ldots, 8$, k along the x axis). The decay of the weights is exponential in k; five to eight components are needed to achieve 1% error (e.g. five for the 0th, seven for the 4th, and eight for the 8th shown in Fig. 6). (Bottom) The first four radial (left) and scale (right) components of the 5th singular function: $r_k^4(\rho)$ and $s_k^4(\sigma)$, $k = 0, \ldots, 3$ (see (16)). (Top right) The real part of one scale-component of a_5, the 5th singular function: $\cos(2\pi\nu_4\theta)s_3^4(\rho)$ (see (17)).

In Fig. 8 reconstructions of the kernel based on a 1% error decomposition are shown for various scales and angles. A maximum of 1% error was imposed on the original steerable decomposition, and again on the scale decomposition of each single a_i. The measured error was 2.5% independent of angle and scale. The total number of filters required to implement a three-octave 1% (nominal, 2.5% real) approximation error of the 3 : 1 Gaussian pair is 16 (rotation) × 8 (scale) = 128. If 10% approximation error is allowed the number of filters decreases by approximately a factor of 4 to 32.

VI. Open Issues and Discussion

A few issues remain open to investigation:

1) The decomposition of the kernels a_i into X-Y-separable ones (Section IV.B) speeds up computations considerably in digital square-lattice implementations. However, the decomposition is done a posteriori, after the SVD decomposition of $F_0(\mathbf{x})$ into $(a_i(\mathbf{x}), b_i(\theta), \sigma_i)$, and, therefore, it may not be the most parsimonious X-Y-separable decomposition. This issue ought to be addressed. One ought to impose X-Y-separability together with steerability and scalability right from the start.
2) Sometimes one would like to generate a discrete decomposition of a family of filters that obeys other constraints than just being the most parsimonious one. For example a) hardware limitations could constrain the shape of the interpolating funtions $b(\theta)$, b) one would like to build pyramid implementations of the decomposition for speeding up the filtering stage (work on this issue has been done by Simoncelli et al. [62]).
3) Another interesting question mentioned in the introduction is the synthesis of the discrete decomposition directly from the specification of an early vision task, rather than passing through the synthesis of a 2D (nD) kernel which then is deformed somewhat arbitrarily. Work in this direction has been done by Hueckel [63], [64], Hummel [40], and Haralick [65] who approached the problem of feature (step edge, line in [64]) detection and localization as one of projecting image neighbourhoods on small-dimension linear subspaces, and deriving the relevant parameters (orientation, position) of the feature from this reduced representation.

   Hummel's approach is particularly interesting: the parameters describing the feature are modelled as continuous random variables. The neighborhood operators ("=" kernels of the linear filters) used to project each neighborhood onto a small-dimensional subspace space are selected using the Karhunen-Loève transform. Such procedure guarantees that the projection maximizes the variance of the parameters and therefore the parameters thus obtained are maximally informative.

   The similarity of the kernels derived by Hueckel and Hummel to the a_i depicted in Fig. 1 is not totally surprising: the polar separability and the fact that the tangential component of the kernels is sinusoidal has to be expected from the fact that one of the parameters in question is a rotation in the plane.

VII. Conclusions

A technique has been presented for implementing families of deformable kernels for early vision applications. A given family of kernels obtained by deforming continuously a tem-

plate kernel is approximated by interpolating a finite discrete set of kernels. The technique may be applied if and only if the family of kernels involved satisfy a compactness condition. This improves upon previous work by Freeman and Adelson on steerable filters in that 1) it is formulated with maximum generality to the case of any compact deformation, or, equivalently any compact family of kernels, and 2) it provides a design technique which is guaranteed to find the most parsimonious discrete approximation.

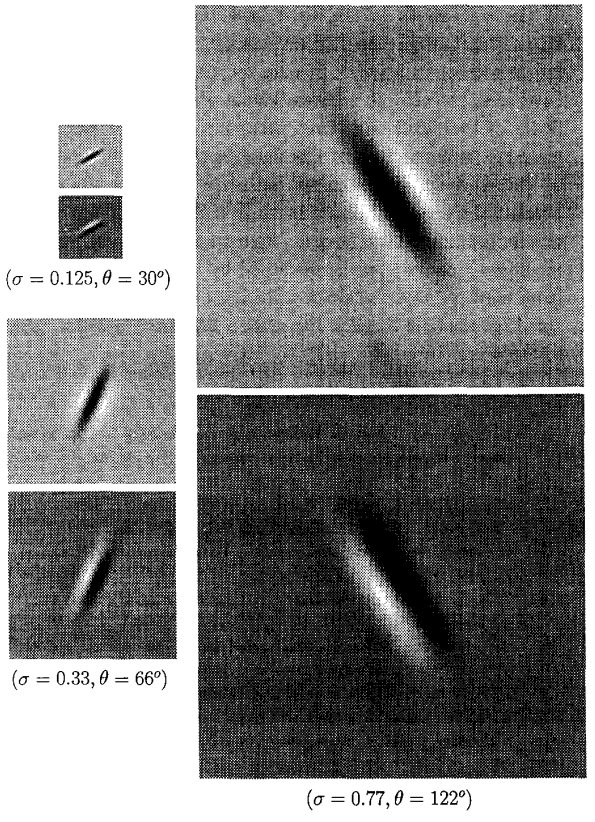

Fig. 8. The kernel at different scales and orientations: the scales are (left to right) 0.125, 0.33, 0.77. The orientations are (left to right) 30°, 66°, 122°. The kernels shown here were obtained from the scale-angle decomposition.

Unlike common techniques used in early vision where the set of orientations is discretized, here the kernel and the response of the corresponding filter may be computed in a continuum for *any* value of the deformation parameters, with no anisotropies. The approximation error is computable a priori and it is constant with respect to the deformation parameter. This allows one, for example, to recover edges with great spatial and angular accuracy.

APPENDIX

A. Proof of Theorem 1

What follows is a derivation to prove Theorem 1. The proof is summarized at the end of the section.

PROOF. THE family of functions F_θ defined by (1) may be thought of as the kernel associated to a linear operator $L : A \to B$, defined by

$$b(\theta) = (La)(\theta) = \int_{\mathbf{R}^2} F_\theta(\mathbf{x}) a(\mathbf{x}) d\mathbf{x} \tag{18}$$

where $A = L^2(\mathbf{R}^2)$, and $B = L^2(\mathbf{S}^1)$, the square integrable functions defined on the plane and the circle, respectively, and $a \in A$, $b \in B$, $\theta \in \mathbf{S}^1$. Let L^* denote the adjoint to L, i.e., the unique linear operator satisfying the equality

$$\langle La, b \rangle_B = \langle a, L^* b \rangle_A \tag{19}$$

with $\langle \cdot, \cdot \rangle_C$ indicating the inner product of a Hilbert space C.

If $\|F\| \leq \infty$ then F_θ is a Hilbert-Schmidt kernel and L has a discrete spectrum; then F_θ can be written as a sum:

$$F_\theta(\mathbf{x}) = \sum_{i=1}^{N} \sigma_i a_i(\mathbf{x}) b_i(\theta) \tag{21}$$

where the σ_i^2 are the nonzero (positive, in fact) eigenvalues of the auto-adjoint operators $LL^* = L \circ L^*$ and $L^*L = L^* \circ L$, and the a_i and b_i are the associated eigenfunctions of L^*L and LL^*, respectively, and N could be infinite. The collection of triples $(\sigma_i, a_i, b_i)_{i=1, \ldots, N}$ is the SVD of L (see e.g., [66]).

Observe that expression (20) is in the desired form of (1), with the additional advantage that the a_i and b_i are orthonormal bases of A and B (see below). Therefore if one could calculate the SVD (i.e., the a_i, b_i, and σ_i) of L explicitly one would be able to write F_θ as in (1), and the problem would be solved.

L^* can be computed from its definition (19) and (18) and is:

$$(L^* b)(\mathbf{x}) = \int_{\mathbf{S}^1} \overline{F_\theta(\mathbf{x})} b(\theta) d\theta \tag{21}$$

hence the LL^* operator is

$$(LL^* b)(\theta) = \int_{\mathbf{S}^1} H(\theta, \theta') b(\theta') d\theta' \tag{22}$$

$$H(\theta, \theta') = \int_{\mathbf{R}^2} F_\theta(\mathbf{x}) \overline{F_{\theta'}(\mathbf{x})} d\mathbf{x} \tag{23}$$

and L^*L is

$$(L^* La)(\mathbf{x}) = \int_{\mathbf{R}^2} K(\mathbf{x}, \mathbf{x}') a(\mathbf{x}') d\mathbf{x}' \tag{24}$$

$$K(\mathbf{x}, \mathbf{x}') = \int_{\mathbf{S}^1} F_\theta(\mathbf{x}) \overline{F_\theta(\mathbf{x}')} d\theta \tag{25}$$

Observe that the kernel associated with LL^* is a function of the difference of its arguments only. To see that change the variable of integration in (23), $y = R_{\theta'}\mathbf{x}$, obtaining $H(\theta, \theta')$

$= H(\theta - \theta', 0) = h(\theta - \theta')$.

The eigenvalue-eigenvector problem for LL^*

$$LL^* b_i = \lambda_i b_i \tag{26}$$

can then be solved explicitly substituting (22) in (26):

$$\int_{S^1} h(\theta - \theta') b_i(\theta') d\theta' = \lambda_i b_i(\theta) \tag{27}$$

The equation holds between the Fourier transforms of the two terms of the shift-invariant convolution of (27):

$$\hat{h}(\nu)\hat{b}_i(\nu) = \lambda_i \hat{b}_i(\nu) \tag{28}$$

The nonzero solutions of this equation are the couples $(\lambda_i, \hat{b}_i)$ s.t. $\lambda_i = \hat{h}(\nu_i)$, and $\hat{b}_i(\nu) = \delta(\nu - \nu_i)$; therefore

$$\sigma_i = \hat{h}(\nu_i)^{1/2} \qquad b_i(\theta) = e^{j2\pi\nu_i\theta} \tag{29}$$

where, by convention, the frequency numbers ν_i are ordered so that $\hat{h}(\nu_i)$ is a nonincreasing sequence: $\hat{h}(\nu_i) \geq \hat{h}(\nu_{i+1}) > 0$. For multiple eigenvalues any linear combinations of the corresponding eigenfunctions (eigenvectors) will also be eigenfunctions.

The eigenfunctions of the L^*L operator cannot be easily determined directly from its integral expression (24) as for LL^*. However one can make use of the fact that $L^* b_i$ is an eigenfunction of L^*L, which can be verified as follows: $L^*L(L^* b_i) = L^*(LL^* b_i) = L^* \lambda_i b_i = \lambda_i (L^* b_i)$. The unit-norm eigenfunctions of L^*L are therefore the functions defined by

$$a_i(\mathbf{x}) = \lambda_i^{-1/2} L^* b_i = \sigma_i^{-1} \int_{S^1} \overline{F_\theta(\mathbf{x})} e^{j2\pi\nu_i\theta} d\theta \tag{30}$$

i.e., at each $\mathbf{x}$, $a_i(\mathbf{x})$ is the conjugate of the Fourier coefficient of $F_\theta(\mathbf{x})$ corresponding to the frequency $-\nu_i$.

In conclusion (the numbers refer to the corresponding statements of the theorem):

1) Follows from the properties of SVD and the fact that the sum (20) is built using the SVD triples.
2) As above.
3) From (30) and the fact that the dimension of the range of L^* is equal to N, the number of nonzero eigenvalues.
4) Follows from SVD properties.
5) Follows from the fact that $\sigma_i \geq \sigma_{i+1}$ and that $\sum_i \sigma_i = \|L\|^2 < \infty$. May be seen directly from SVD properties.
6) From (13) one can see that $F_\theta^{[n]}$ is a function of $|\mathbf{x}|$ and $\theta - \phi(\mathbf{x})$ only.
7) The functions a_i are linearly independent, so any collection of n of them spans an n-dimensional subspace of $L^2(\mathbf{R}^2)$. This is the same subspace spanned by any linearly independent collection $F_{\theta_i}^{[n]} i = 1, \ldots, n$. The thesis follows from the fact $F_{\theta_i}^{[n]} = F_{\theta=0}^{[n]} \circ R_{\theta_i}$. The coefficients α_i can be obtained with the usual change of basis formulae. □

ACKNOWLEDGEMENTS

I am very grateful to Massimo Porrati, Alberto Grunbaum, David Donoho, Federico Girosi, Mohammad Naraghi, and Frank Ade for giving me good advice and useful references. I would also like to acknowledge useful conversations with Ted Adelson, Josef Bigün, Stefano Casadei, Rachid Deriche, Charles Desoer, Peter Falb, Bill Freeman, Milan Jovovic, Takis Konstantopoulos, Olaf Kübler, Paul Kube, Reiner Lenz, Jitendra Malik, Stephane Mallat, Sanjoy Mitter, Richard Murray, Shankar Sastry, and Doug Shy. The comments of one anonymous reviewer are also gratefully acknowledged.

The simulations in this work have been carried out using Paul Kube's "viz" image-manipulation package, and would have taken much longer without Paul's very friendly support. The images in this paper have been printed with software kindly provided by Eero Simoncelli. Some of the simulations have been run on a computer generously loaned to the author by Prof. Canali of the Università di Padova.

This work was in part carried out while the author was with the International Computer Science Institute at Berkeley and at the Laboratory for Information and Decision Systems of M.I.T. with the Center for Intelligent Control Systems sponsored by U.S. Army Research Office grant number DAAL 03-86-K-0171. This research was in part sponsored by National Science Foundation Research Initiation grant IRI 9211651.

REFERENCES

[1] D. Hubel and T. Wiesel, "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex.," *J. Physiol. (Lond.)*, vol. 160, pp 106–154, 1962.

[2] D. Hubel and T. Wiesel, "Receptive fields of single neurones in the cat's striate cortex.," *J. Physiol. (Lond.)*, vol. 148, pp. 574–591, 1959.

[3] B. Horn, "The binford-horn linefinder," Technical report, MIT AI Lab. Memo 285, 1971.

[4] T. Binford, "Inferring surfaces from images," *Artificial Intelligence*, vol. 17 pp. 205–244, 1981.

[5] G.H. Granlund, "In search of a general picture processing operator," *Computer Graphics and Image Processing*, vol. 8, pp. 155–173, 1978.

[6] H. Knuttson and G.H. Granlund, "Texture analysis using two-dimensional quadrature filters," *Workshop on Computer Architecture for Pattern Analysis and Image Database Management*, pp. 206–213, IEEE Computer Society 1983.

[7] M. Kass, "Computing visual correspondence," *Proc. Image Understanding Workshop*, (McLean, Va.), pp. 54–60, Science Applications, Inc., June 1983.

[8] P. Burt and E. Adelson, "The laplacian pyramid as a compact image code," *IEEE Trans. Communications*, vol. 31, pp. 532–540, 1983.

[9] E. Adelson and J. Bergen, "Spatiotemporal energy models for the perception of motion," *J. Opt. Soc. Am.*, vol. 2, no. 2, pp. 284–299, 1985.

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

[11] D. Heeger, "Optical flow from spatiotemporal filters," *Proc. First Int'l Conf. Computer Vision*, pp. 181–190, 1987.

[12] M. Turner, "Texture discrimination by gabor functions," *Biol. Cybern.*, vol. 55, pp. 71–82, 1986.

[13] M. Morrone and R. Owens, "Feature detection from local energy," *Pattern Recognition Letters*, vol. 6, pp. 303–313, 1987.

[14] M. Morrone and D. Burr, "A model of human feature detection based on matched filters," *Robots and Biological Systems*, P. Dario and G. Sandini, ed., Academic Press, 1990.

[15] P. Parent and S. Zucker, "Trace inference, curvature consistency, and curve detection," *IEEE Trans. Pattern Analysis and . Machine Intelligence*, vol. 11, no. 8, pp. 823–839, 1989.

[16] I. Fogel and D. Sagi, "Gabor filters as texture discriminators," *Biol. Cybern.*, vol. 61, pp. 103–113, 1989.

[17] J. Malik and P. Perona, "Preattentive texture discrimination with early vision mechanisms," *J. Opt. Soc. Am. A*, vol. 7, no. 5, pp. 923–932, 1990. Also ap-

peared as a chapter in *Computer vision: advances and applications*, R. Kasturi and R. Jain, ed., IEEE Computer Society Press, 1991.

[18] A. Bovik, M. Clark, and W. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 1, pp. 55–73, 1990.

[19] S. Zhong and S. Mallat, "Compact image representation from multiscale edges," *Proc. Third Int'l Conf. Computer Vision*, IEEE Computer Society, Osaka, 1990.

[20] D. Jones and J. Malik, "Computational stereopsis–Beyond zero-crossings," *Invest. Ophtalmol. Vis. Sci. (Supplement)*, vol. 31, no. 4, p. 529, 1990.

[21] D. Jones and J. Malik, "Using orientation and spatial frequency disparities to recover 3d surface shape—A computational model," *Invest. Ophtalmol. Vis. Sci. (Supplement)*, vol. 32, no. 4, p. 710, 1991.

[22] W. Freeman and E. Adelson, "Steerable filters for early vision, image analysis and wavelet decomposition," *Third Int'l Conf. Computer Vision*, pp. 406–415, IEEE Computer Society, 1990.

[23] P. Perona and J. Malik, "Detecting and localizing edges composed of steps, peaks and roofs," *Proc. Third Int'l Conf. Computer Vision*, pp. 52–57, IEEE Computer Society, Osaka, 1990.

[24] E. Adelson and J. Bergen, "The plenoptic function and the elements of early vision," *Computational Models of Visual Processing*, M. Landy and J. Movshon, ed., MIT Press, 1991. Also appeared as MIT-MediaLab-TR148. Sept. 1990.

[25] J. Koenderink and A. van Doorn, "Representation of local geometry in the visual system," *Biol. Cybern.*, vol. 55, pp. 367–375, 1987.

[26] J. Koenderink and A. van Doorn, "Receptive field families," *Biol. Cybern.*, vol. 63, pp. 291–298, 1990.

[27] R. Lenz, "Group theoretical methods in image processing," *Lecture Notes in Computer Science*, vol. 413, Springer Verlag, 1990.

[28] B. ter Haar Romeny, L. Florack, J. Koenderink, and M. Viergever, "Scale-space, its natural operators and differential invariants," *Int'l Conf. Information Processing in Medical Imaging*, LCNS series vol. 511, pp. 239–255, Springer Verlag, 1991.

[29] J. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1160–1169, 1985.

[30] J. Daugman, "Two-dimensional spectral analysis of cortical receptive field profiles," *Vis. Research*, vol. 20, pp. 847–856, 1980.

[31] R. Young, "The gaussian derivative theory of spatial vision: Analysis of cortical cell receptive field line-weighting profiles," Technical Report GMR-4920, General Motors Research, 1985.

[32] M. Morrone and D. Burr, "Feature detection in human vision: a phase dependent energy model," *Proc. R. Soc. Lond. B*, vol. 235, pp. 221–245, 1988.

[33] D. Jones and J. Malik, "A computational framework for determining stereo correspondence from a set of linear spatial filters," *Proc. Second European Conf. Computer Vision*, G. Sandini, ed., LNCS-Series Vol. 588, Springer-Verlag, pp. 395–410, 1992.

[34] D. Jones and J. Malik, "Determining 3d shape from orientation and spatial frequncy disparities ii—Using correspondence image patches," UCB-CSD 91-657, U.C.Berkeley—CS Dept., Oct, 1991.

[35] J. Malik and Z. Gigus, "A model for curvilinear segregation," *Invest. Ophtalmol. Vis. Sci. (Supplement)*, vol. 32, no. 4, p. 715, 1991.

[36] L. Rosenthaler, F. Heitger, R. von der Heydt, and O. Kübler, "Detecting general edges and keypoints," Tech. Rep. 130, IKT/Image science lab ETH-Zuerich, 1991.

[37] F. Heitger, L. Rosenthaler, R. von der Heydt, E. Peterhans, and O. Kübler, "Simulation of neural contour mechanismn: From single to end-stopped cell," *Vis. Research*, vol. 32, no. 5, pp. 963–981, 1992.

[38] J. Daugman, "Complete discrete 2d gabor transform by neural networks for image analysis and compression," *IEEE Trans. Acoustics, Speech, Signal Processing*, vol. 36, no. 7, pp. 1169–1179, 1988.

[39] M. Morrone, D. Burr, J. Ross, and R. Owens, "Mach bands depend on spatial phase," *Nature*, vol. 324, no. 6094, pp. 250–253, 1986.

[40] R. Hummel, "Feature detection using basis functions," *Comp. Vision, Graphics and Image Proc.*, vol. 9, pp. 40–55, 1979.

[41] H. Knuttson, *Filtering and reconstruction in image processing*. PhD thesis, Linköping Inst. of Technology, 1982.

[42] J. Bigün, G. Granlund, and J. Wiklund, "Multidimensional orientation estimation with applications to texture analysis and optical flow," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 8, pp. 775–790, 1991.

[43] W. Freeman and E. Adelson, "Steerable filters," *Topical Meeting on Image Understanding and Machine Vision*, vol. 14, Optical Soc. America, technical digest series, Cape Cod, June 1989.

[44] W. Freeman and E. Adelson, "Steerable filters for image analysis," Technical Report 126, MIT, Media Laboratory, 1990.

[45] W. Freeman and E. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, pp. 891–906, 1991.

[46] P. Perona and J. Malik, "Detecting and localizing edges composed of steps, peaks and roofs," Tech. Rep. UCB/CSD 90/590, Computer Science Division (EECS), U.C.Berkeley, 1990.

[47] P. Perona and J. Malik, "Boundary detection using quadratic filters: performance criteria and experimental assessment.," *Proc. SPIE Conf. Applications of Artificial Intelligence*, vol. 1708, Orlando, April 1992.

[48] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. Information Theory*, vol. 38, no. 2, pp. 587–607, 1992.

[49] P. Perona, "Finite representation of deformable functions," Tech. Rep. 90-034, Int'l Computer Science Inst., Berkeley, Calif., 1990.

[50] P. Perona, "Deformable kernels for early vision," *Proc. IEEE CS Conf. Computer Vision and Pattern Recognition*, pp. 222–227, Maui, June 1991.

[51] P. Perona, "Deformable kernels for early vision," Tech. Rep. 2039, LIDS-MIT, Oct. 1991.

[52] P. Perona, "Steerable-scalable kernels for edge detection and junction analysis," *Proc. Second European Conf. Computer Vision*, G. Sandini, ed., LNCS-Series Vol. 588, Springer-Verlag, pp. 3–18, 1992. Reproduced in Image and Vision Computing, vol 10, pp. 663-672, 1992.

[53] G. Choquet, *Lectures on Analysis*, vol. I. New York: W. A. Benjamin Inc., 1969.

[54] J. Dieudonne, *Foundations of Modern Analysis*. New York: Academic Press, 1969.

[55] A. Pinkus, *n-Widths in Approximation Theory*. Springer Verlag, 1985.

[56] S. Treitel and J. Shanks, "The design of multistage separable planar filters," *IEEE Trans. Geosci. Electron.*, vol. 9, no. 1, pp. 10–27, 1971.

[57] W.K. Pratt, "An intelligent image processing display terminal," *Proc. SPIE Tech. Symp.*, vol 27, San Diego, 1979.

[58] R. Deriche and J.-F. Abramatic., "Design of 2d recursive filters using singular value decomposition techniques," *Proc. ICASSP*, 1982.

[59] J. Dongarra, C. Moler, J. Bunch, and G. Stuart, *Linpack, User's Guide*. Soc. Industrial and Applied Mathematics, 1972.

[60] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes in C*. Cambridge University Press, 1988.

[61] J.-F. Abramatic and O. Faugeras, "Sequential convolution techniques for image filtering," *IEEE Trans. Acoustics, Speech, Signal Processing*, vol. 30, no. 1, pp. 1–10, 1982.

[62] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. Information Theory*, vol.38, pp. 587-607.

[63] M. Hueckel, "An operator which locates edges in digitized pictures," *J. Assoc. Comp. Mach.*, vol. 18, no. 1, pp. 113–125, 1971.

[64] M. Hueckel, "A local visual operator which recognizes edges and lines," *J. Assoc. Comp. Mach.*, vol. 20, no. 4, pp. 643–647, 1973.

[65] R. Haralik, "Digital step edges from zero crossing of second directional derivatives," *IEEE Trans. Pattern Analysis and Machine Intelligence.*, vol. 6, no. 1, pp. 58–68, 1984.

[66] W. Rudin, *Functional Analysis*. McGraw-Hill, 1973.

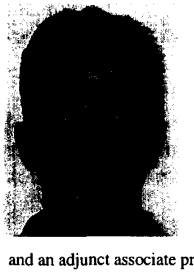

Pietro Perona received the DEng degree in electrical engineering from the University of Padova in 1985 and the PhD degree in electrical engineering and computer science from the University of California at Berkeley in 1990.

Dr. Perona was a postdoctoral fellow with the International Computer Science Institute at Berkeley and the Laboratory of Decision and Information Systems at M.I.T., and a *recercatore* with the University of Padova from 1990 to 1991. He has been an assistant professor with the California Institute of Technology since 1991 and an adjunct associate professor with the University of Padova since 1993.

Dr. Perona's research interests include computational vision, visual psychopsysics, and the modeling of human vision.