

A continuous Formulation of intrinsic Dimension

Norbert Krüger^{1,2} Michael Felsberg³

Computer Science and Engineering¹
Aalborg University Esbjerg
6700 Esbjerg, Denmark
nk@cs.aue.auc.dk

Electrical Engineering³
Linköping University
SE-58183 Linköping, Sweden
mfe@isy.liu.se

Abstract

The intrinsic dimension (see, e.g., [29, 11]) has proven to be a suitable descriptor to distinguish between different kind of image structures such as edges, junctions or homogeneous image patches. In this paper, we will show that the intrinsic dimension is spanned by two axes: one axis represents the variance of the spectral energy and one represents the a weighted variance in orientation. Moreover, we will show in section *that the topological structure of intrinsic dimension has the form of a triangle*. We will review diverse definitions of intrinsic dimension and we will show that they can be subsumed within the above mentioned scheme. We will then give a concrete continuous definition of intrinsic dimension that realizes its triangular structure.

1 Introduction

Natural images are dominated by specific local sub-structures, such as edges, junctions, or texture. Sub-domains of Computer Vision have analyzed these sub-structures by making use of certain concepts (such as, e.g., orientation, position, or texture gradient). These concepts were then utilized for a variety of tasks, such as, edge detection (see, e.g., [6]), junction classification (see, e.g., [27]), and texture interpretation (see, e.g., [26]). However, before interpreting image patches by such concepts we want know whether and how these apply. For example, the idea of orientation does make sense for edges or lines but not for a junction or most textures. As another example, the concept of position is different for a junction compared to an edge or an homogeneous image patch. For a junction the position can be unambiguously defined by the point of intersection of lines, for edges the aperture problem leads to a definition of the position as a one-dimensional manifold and for an homogeneous image patch it is impossible to define a position in terms of local signal attributes. Hence, before we apply concepts like orientation or position, we want to classify image patches according to their junction-ness, edge-ness or homogeneous-ness.

The intrinsic dimension (see, e.g., [29, 11]) has proven to be a suitable descriptor in this context. Homogeneous image patches have an intrinsic dimension of zero (i0D), edge-like structures are intrinsically 1-dimensional (i1D) while junctions and most textures have an intrinsic dimension of two (i2D). There exists also related classifications such as the rank of a image patch [17], the rank taking discrete values zero, one, or two. Another related formulation is the distinction between constant, simple and isotropic signals [19]. The association of intrinsic dimension to a local image structure has mostly be

²This work was conducted at the departement of Psychology, University of Stirling, FK94LA Stirling, Scotland.

done by a discrete classification [29, 11, 19]. To our knowledge, so far there exists no continuous definition of intrinsic dimensionality that covers all three possible cases (i0D, i1D, and i2D). However, there exist attempts to find a continuous formulation between i1D and i2D signals [17].

In contrast to, e.g. curvature estimators (see, e.g., [2, 25]), the intrinsic dimensionality does not make any assumption about specific structural attributes of the signal but is based on a purely statistical criterion: The concept of curvature does make sense for curved lines but not for junctions or most complex textures. However, the intrinsic dimension is a sensible descriptor also for these kind of signals (see also [20]).

In section 2.1, we will show that the intrinsic dimension is a local descriptor that is spanned by two axes: one axis represents the variance of the spectral energy and one represents the a weighted variance in orientation. In this paper, we will review diverse definitions of intrinsic dimension. In section 2.2, we will show that they can be subsumed within the above mentioned scheme. Since the intrinsic dimension is a two-dimensional structure, no continuous one-dimensional definition is sensible. Moreover, we will show in section 2.1 *that the topological structure of intrinsic dimension essentially has the form of a triangle*. We will then give one possible concrete definition of intrinsic dimension that realizes its triangular structure in section 3.1.

A classification of edge-ness or corner-ness based on a local image patch without taking the context into account always faces the problem of the high degree of ambiguity of visual information (see, e.g., [1]). Taking into account this ambiguity we do not want to come to a final decision about the junction-ness of edge-ness of an image patch but we want to associate confidences to such classifications. Assigning confidences instead of binary decisions at low level stages of processing has been proven useful since it allows for stabilizing such local classifications according to the context (see, e.g., [1, 21]). By making use of barycentric coordinates (see, e.g., [7]), we will utilize the triangular structure of intrinsic dimension to express confidences for the different possible interpretation in section 3.2. This leads to *continuous definition of intrinsic dimensionality that covers i0D, i1D and i2D signals*. Finally, in section 4 we show examples of our continuous classification of image patches of different intrinsic dimension.

To our knowledge, this paper is the first work that makes the triangular structure of intrinsic dimensionality explicit and which gives a continuous definition that covers all three possible cases of intrinsic dimension.

2 The Concept of intrinsic Dimensionality

The *intrinsic dimensionality* in image processing is a formalization of what is commonly called "edgeness" vs. "junction-ness". The term intrinsic dimensionality itself is much more general. In [4], p. 314, it says that "a data set in d dimensions is said to have an *intrinsic dimensionality* equal to d' if the data lies entirely within a d' -dimensional subspace", but indeed, the concept of intrinsic dimensionality is much older [28].

In image processing, the intrinsic dimensionality was introduced by [29] to define heuristically a discrete distinction between edge-like and corner-like structures. However, here we want to adopt the more general definition in [4] to image processing. For this, we have to consider the *spectrum* of an image patch (see figure 1):

- if the spectrum is concentrated in a point¹, the image patch has an intrinsic dimensionality of null (i0D),

¹Note that due to the Hermitian spectrum of a (real valued) image, this point can only be the origin, i.e., the DC component.

- if the spectrum is concentrated in a line², the image patch has an intrinsic dimensionality of one (i1D), and
- otherwise the image patch has an intrinsic dimensionality of two (i2D).




Figure 1: Illustration intrinsic dimensionality. In the image on the left, three neighborhoods with different intrinsic dimensionalities are indicated. The other three images show the local spectra of these neighborhoods, from left to right: i0D, i1D, and i2D.

Each of these three cases can be characterized more vividly. Constant image patches correspond to i0D patches. Edges, lines, and sinusoid-like textures obtained by projecting 1D functions (simple signals [17]) correspond to i1D patches. All other structures like corners, junctions, complex textures, and noise correspond to i2D patches.

Taking a closer look at the concept of intrinsic dimensionality, two fundamental problems pop up:

1. The intrinsic dimensionality as it is defined above is a discrete feature in $\{i0D, i1D, i2D\}$. However, every real signal consists of a combination of intrinsic dimensionalities – there are hardly any totally constant or ideal i1D image patches in real images. Hence, we would like to have a *continuous* definition of intrinsic dimensionality.
2. The topology of the iD-space is yet undefined. In case of a discrete space, the relations between the different intrinsic dimensionalities is obvious, all dimensionalities are mutually adjacent. The topology of the continuous iD-space is considered in the subsequent section.

In the following section we discuss a new model for representing the intrinsic dimensionality in a continuous, topologically appropriate way. The subsequent section gives an overview of known methods for estimating the intrinsic dimensionality and relates them to our new model.

2.1 The Intrinsic Dimensionality has a 2D Triangular Structure

For the estimation of the intrinsic dimensionality of an image patch, we need to apply a measure for the spread of the spectral data, either to a point or to a line. The classical approach from statistics for such a measure is the *variance* of the data. Since a change of the coordinate system results in new stochastic variables, the computation of the variance depends on the coordinate system, for instance in Cartesian coordinates vs. polar coordinates. Different coordinate systems lead to further diversification of practical approaches.

²With the same argument as in footnote 1, this line goes through the origin.

To be more concrete, the variance of the spectral data with respect to the origin in a cartesian coordinate system is defined by

$$\sigma_O^2 = \frac{1}{N} \iint_{\Omega} |\mathbf{u}|^2 |F(\mathbf{u})|^2 d\mathbf{u} , \quad (1)$$

where \mathbf{u} is the frequency vector, Ω is the region of integration in the Fourier domain³ and

$$N = \iint_{\Omega} |F(\mathbf{u})|^2 d\mathbf{u} \quad (2)$$

is a normalization constant. The variance with respect to a line is given by

$$\sigma_L^2 = \min_{\mathbf{n}} \frac{1}{N} \iint_{\Omega} |\mathbf{n}^T \mathbf{u}|^2 |F(\mathbf{u})|^2 d\mathbf{u} , \quad (3)$$

where \mathbf{n} is obtained to be parallel to i1D signals, i.e., it represents the orientation. The variance σ_O^2 defines some kind of measure of the local grey level variation whereas the the variances σ_L^2 reflects the dynamic perpendicular to the main orientation.

If we change to polar coordinates $\mathbf{u} \mapsto (q, \theta)$, we get two new variances, the radial variance

$$\sigma_R^2 = \frac{1}{N'} \int_0^Q q^2 \int_0^{2\pi} |F(q \cos \theta, q \sin \theta)|^2 d\theta dq , \quad (4)$$

where Q is the radius of Ω , and the angular variance

$$\sigma_A^2 = \min_{\theta_0} \frac{1}{N'} \int_{\theta_0-\pi}^{\theta_0+\pi} (\theta - \theta_0)^2 \int_0^Q |F(q \cos(\theta - \theta_0), q \sin(\theta - \theta_0))|^2 dq d\theta , \quad (5)$$

where the normalization constant N' is given similar to (2), performing the integration in polar coordinates. The angle θ_0 represents the local orientation.

The two characterizations (σ_O^2, σ_L^2) and (σ_R^2, σ_A^2) are different in detail, but related. The most important difference between the two variances σ_O^2 and σ_R^2 is the different weighting of the frequency components due to the missing Jacobian of the coordinate transform. The two variances σ_L^2 and σ_A^2 differ more essentially, since σ_A^2 becomes undefined for $\sigma_R^2 = 0$. The orientation variance σ_A^2 corresponds to formulations of *intensity invariant measures of the 'i1D-ness'*. This intensity invariance prevents a probabilistic, triangular formulation of intrinsic dimensionality since the i0D case is neglected. Examples for such traditional, intensity invariant measures are the coherence [19] (see also next section) and the isotropy factor [11].

For the idealized cases of purely i0D, i1D, and i2D signals, the variances obviously behave as given in table 1. By a proper normalization ($\tilde{\sigma}_O^2 = k\sigma_O^2$, etc.), the entries "large" in this table can be replaced by "1", yielding an overall range of $[0, 1] \times [0, 1]$, i.e., the iD-space spanned by $\tilde{\sigma}_R^2$ and $\tilde{\sigma}_A^2$ corresponds to a *2D square*. The entry "undefined", however, cannot be simply replaced by a value between zero and one, all values coexist with the same right. In other words, one edge of the square is singular. To solve for this singular edge, a straightforward idea is to multiply $\tilde{\sigma}_A^2$ with $\tilde{\sigma}_R^2$ (see section 3), which can be considered as a replacement for the Jacobian, i.e., we consider a space similar to $(\tilde{\sigma}_O^2, \tilde{\sigma}_L^2)$ instead. Since the $(\tilde{\sigma}_R^2, \tilde{\sigma}_A^2)$ space is a 2D square, we obtain a *2D triangle* for $(\tilde{\sigma}_O^2, \tilde{\sigma}_L^2)$, see figure 2.

³In practice, this is mostly a Gaussian window, i.e., we consider the windowed Fourier transform (2D version of the short-time Fourier transform).

intrinsic dimensionality	i0D	i1D	i2D
σ_O^2	0	large	large
σ_L^2	0	0	large
σ_R^2	0	large	large
σ_A^2	undefined	0	large

Table 1: Intrinsic dimensionality and variances. A zero variance means a ideal concentration of the spectral data.




Figure 2: About the topology of iD-space. Left: traditional iD-space (square), center: our iD-space (triangle), right: parametrization of the iD-triangle by barycentric coordinates.

Each of the corners of the triangle corresponds to a certain intrinsic dimensionality. The topology of the triangle allows to vary the intrinsic dimensionality continuously from any case to any other case. This observation is very important in practice since, as stated further above, every real signal consists of a combination of intrinsic dimensionalities. The parameterization of the iD-triangle is described in detail in section 3.

2.2 Approaches for Estimating the Intrinsic Dimensionality

The various approaches which occurred in the literature so far mainly differ with respect to two aspects: (1) the computation of the variances and (2) the coordinate system. Nearly all systematic approaches to measure the intrinsic dimensionality are known as or are equivalent to the *structure tensor* [3, 16].

Basically, the variances can either be computed by outer products of first order derivatives or by combinations of quadrature filter responses, see [17, 19] for an overview. There are other but still related methods, e.g., polynomial expansions [10] and higher order spherical harmonics [13]. Most approaches make use of Cartesian coordinates to compute and to represent the variances, but an evaluation in polar coordinates is at least a plausible alternative, see section 3 and [8].

The first approach to what is nowadays called structure tensor is based on averaging the outer product of derivatives. This method was independently invented by Bigün and Granlund [3] and Förstner and Gülch [16]. In [15] a deeper analysis of the structure tensor from a statistical point of view is developed. The idea is to approximate the auto-covariance function by a truncated Taylor series expansion in the origin. The term which is obtained by this expansion is given by

$$\mathbf{J} = \int_{\Omega} \mathbf{u} \mathbf{u}^T |F(\mathbf{u})|^2 d\mathbf{u} . \quad (6)$$

Applying the power theorem [5] and the derivative theorem, we end up with

$$\mathbf{J} = \int_{\omega} (\nabla f(\mathbf{x})) (\nabla f(\mathbf{x}))^T d\mathbf{x} , \quad (7)$$

where ω is a local region of integration in the spatial domain (a Gaussian window in case of a windowed Fourier transform). This tensor \mathbf{J} can be visualized by an ellipse, where the length of the two main axes correspond to the two eigenvalues λ_1 and λ_2 of the tensor. The mean of the two eigenvalues (the trace of the tensor) corresponds to the variance with respect to the origin σ_O^2 , and the smaller eigenvalue λ_2 corresponds to the line-variance σ_L^2 . Therefore, the two axes of the iD-triangle are given by $(\lambda_1 + \lambda_2)/2$ and λ_2 and an appropriate normalization.

The tensor feature which is typically used in context of estimating the intrinsic dimensionality is the *coherence*⁴ [3]:

$$c = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} . \quad (8)$$

The coherence c is related to the variances σ_O^2 and σ_L^2 (and to $\tilde{\sigma}_A^2$) by $c = 1 - 2\sigma_L^2/\sigma_O^2 \approx 1 - \tilde{\sigma}_A^2$. A common method for distinguishing i1D and i2D structures is to threshold the coherence. This is also the theoretic background of the Harris-Stephens corner detector [18]. A drawback of all coherence-based methods is *that an additional energy threshold has to be applied in order to single out constant (i1D) regions*. Our new triangle model (figure 2, section 3) allows us to postprocess the iD information without applying threshold at this early step.

A method which is related to the structure tensor but which is different in detail is based on generalized 2D quadrature filters [13]. The idea of this approach is to compute responses of steerable quadrature filters which are adapted to the main orientation of the signal and to the perpendicular orientation. The filters are polar separable and window out the information in the respectively perpendicular orientation. The effective amplitude response of the filter set is isotropic.

The resulting feature vector consists of five features among which we find the local orientation, the local amplitude with respect to the local orientation A_M , and the local amplitude perpendicular to the local orientation A_m . Local amplitudes computed by a quadrature filter are related to variances in the Fourier domain. Assuming that a quadrature filter has a sufficiently small bandwidth, the filter output approximates the Fourier components at the center frequency [17], page 171. Hence, the local amplitude increases with increasing variance of the spectrum.

The squareroot of the ratio of the two amplitudes $c_I = \sqrt{A_m/A_M}$ is called *isotropy factor* and it corresponds to $\sqrt{(1-c)/(1+c)}$ [11]. Hence, the orientation variance $\tilde{\sigma}_A^2$ is given by

$$\tilde{\sigma}_A^2 \approx 2 \frac{\sigma_L^2}{\sigma_O^2} = 1 - c = \frac{2c_I^2}{1 + c_I^2} = \frac{2A_m}{A_M + A_m} .$$

Due to the isotropy of the filter set, the mean of the two amplitudes corresponds to a total local amplitude of the signal and hence to the variance σ_O^2 . Therefore, after normalization the amplitudes can be used for parameterizing the iD-triangle: $(A_M + A_m)/2$ as the first coordinate and A_m as the second coordinate. Indeed, evaluating A_m and applying a threshold has been used for corner detection in [13].

3 Triangular Definition of intrinsic Dimension

Having shown in section 2.1 that the topological structure of intrinsic dimensionality is essentially a triangle, we now derive a realization of intrinsic dimensionality that makes use of its triangular structure. Instead of a binary classification (as done, e.g., in [29, 11,

⁴In [16] the coherence is squared, which is unnecessary if the eigenvalues are ordered.

19)), we compute 3 values $c_{0D}, c_{1D}, c_{2D}, c_i \in [0, 1]$ that code confidences for the intrinsic 0–dimensionality, intrinsic 1–dimensionality and intrinsic 2–dimensionality of the signal.

In section 3.1 we will concretize the origin–variance and line–variance introduced in section 2.1 and use these measures to span a triangle whose corners represent the extremes of purely i0D, i1D and i2D signals. In section 3.2 we then use barycentric coordinates to assign the confidences.

3.1 Local Amplitude and Orientation Variance as two axes spanning the Triangle

Our image processing starts with a filter operation which is based on generalized quadrature filters [14]. These filters perform a *split of identity*, i.e., the signal becomes orthogonally divided into its amplitude m (indicating the likelihood of the presence of a structure), its geometric information (orientation) θ and its phase φ .

We express our realization of the intrinsic dimensionality triangle in polar coordinates. To compute the origin–variance we first apply a normalization function N that transfers the amplitude m that has values in $[0, \infty]$ to the interval $[0, 1]$ by performing a smooth thresholding using a sigmoidal function. The shape of the sigmoid function does depend on the local and global contrast. In this way even at low contrast image patches image structures can be detected.⁵ Assuming a sufficiently small bandwidth of our filters, our measure for the origin–variance at a pixel position \mathbf{x}_0 is simply given by the normalized amplitude (see section 2.2):

$$\hat{\sigma}_R = N(m(\mathbf{x}_0)).$$

To compute our measure for line–variance at pixel position \mathbf{x}_0 , we compute a weighted variance measure of the local orientation. First, we define a set $A(\mathbf{x}_0)$ representing the local neighbourhood of \mathbf{x}_0 and we compute the mean orientation $E_A[\theta]$ on A . Weighting is performed according to the normalized magnitude. Our measure for line variance then becomes

$$\hat{\sigma}_L = \tilde{\sigma}_A^2 \cdot \tilde{\sigma}_R^2 = \sum_{\mathbf{x} \in A} (N(m(\mathbf{x})) \cdot d(\theta(\mathbf{x}), E_A[\theta])).$$

Note that $\sum_{\mathbf{x} \in A} d(\theta(\mathbf{x}), E_A[\theta])$ basically represent $\tilde{\sigma}_A^2$ and the multiplication with $N(m(\mathbf{x}))$ corresponds to the multiplication with $\tilde{\sigma}_R^2$.

The metric d takes the singularity of the orientation at 0 and π into account and performs a normalisation that ensures that $\hat{\sigma}_L$ takes values in $[0, 1]$. The measure $\hat{\sigma}_L$ defines the second axis of our triangle.

As a final step we apply the squashing function $f(x) = x^c$ to steer the distribution of values in $[0, 1]$. Origin–variance and line variance are finally defined by

$$\begin{aligned} \sigma_O &= \hat{\sigma}_0^{c_1} \\ \sigma_L &= \hat{\sigma}_L^{c_2} \end{aligned}$$

where the parameters $c_1 = \frac{1}{6}$ and $c_2 = \frac{1}{2}$ have to be proven useful. σ_O and σ_L span the triangle (see figure 2). Note that by definition it holds $\sigma_L < \sigma_O$.

Since we have defined the axes of our triangle we can now associate the different intrinsic dimensions to its corners:

An intrinsically zero dimensional (i0D) image patch is characterized by a low origin variance ($\sigma_O \approx 0$). Then it also holds $\sigma_L \approx 0$ since $\sigma_L < \sigma_O$ by definition. In the triangle

⁵This normalization has been proven to be useful in the object recognition system [22] where it is discussed in detail.

shown in figure 2 (right) intrinsically zero dimensional (i0D) image patches correspond to the coordinate $(0, 0)$. Although $m \approx 0$, the local image patch can also be a projection of a 3D–edge (that usually corresponds to i1D signals) or a junction (that usually corresponds to i2D signals). The low contrast may be caused by e.g., accidental background–object constellation or an accidental surface/illumination constellation. To account for these ambiguities we will (based on the representation introduced here) define confidences that express the likelihood of the signal being i0D, i1D or i2D.

An intrinsically one dimensional image patch is characterized by a high origin variance and a low line variance within the image patch. In the triangle in figure 2 (right) this corresponds to the coordinate $(1, 0)$. Note that orientation can only be meaningfully associated to an intrinsically one–dimensional signal patch. In contrast, for a homogenous image patch (i0D) or a junction (i2D) the concept of orientation does not make any sense. With an intrinsically one–dimensional image patch specific problems are associated, for example the aperture problem which is less severe (or non existent) for intrinsically two–dimensional signals.

An intrinsically two dimensional image patch is characterized by high origin variance and high line variance. This corresponds to the coordinate $(1, 1)$ in the triangle shown in figure 2 (right). A parametric description of 2D–image patches is more difficult since there are at least two possible 3D–sources for an intrinsically two–dimensional image patch. First, it may be caused by edges meeting in a point or it may be caused by texture. The underlying 3D–description would be different. A texture is most likely produced by a surface–like structure while a junction most likely is associated to a specific 3D–depth discontinuity.

3.2 Coding intrinsic dimensionality by barycentric coordinates:

Having defined a triangle with its corners representing the extremes in intrinsic dimensionality, we can now code confidences associated to the different intrinsic dimensions (c_{0D}, c_{1D}, c_{2D}) by using barycentric coordinates (see, e.g., [7]). Given a point inside a triangle, the Barycentric coordinates describe twice the area of the triangle opposite to the corners of triangle (see figure 2).

A measurement of σ_0 and σ_L^2 defines a point inside the triangle $(0, 0)$, $(0, 1)$, $(1, 1)$:

$$\mathbf{p} = (p_x, p_y) = (\sigma_0, \sigma_L).$$

Our confidences are the barycentric coordinates of this point:

$$\begin{aligned} c_{0D} &= 1 - p_x \\ c_{1D} &= p_x - p_y \\ c_{2D} &= p_y \end{aligned}$$

Note that since $0 \leq p_y \leq p_x \leq 1$ and $p_x \in [0, 1]$ it holds $0 \leq c_i \leq 1$. The three confidences add up to one since

$$c_{0D} + c_{1D} + c_{2D} = (1 - p_x) + (p_x - p_y) + p_y = 1.$$

4 Simulations

We have applied our definition of intrinsic dimension within a new kind of image representation which is based on multi–modal Primitives (see, e.g. [24]). These Primitives carry information about orientation, colour, optic flow, depth in a condensed way and are used for scene analysis in the European project ECOVISION [9]. To all attributes in the




Figure 3: Primitives of different intrinsic dimensionality (i2D signals are indicated by a star and i1D signals by a line at its centers, i0D signals have no special indicator but have smaller radius. For some Primitives the triangular representation is shown.

different modalities confidences are associated that are subject to contextual modification. Our continuous definition of intrinsic dimension is used as an additional descriptor that codes information about the edge-ness or junction-ness of the Primitive. This allows for, e.g., a use of orientation information for i1D structures only. Figure 3 shows the extracted Primitives from an image and for some of them the position in the triangular representation of the intrinsic dimensionality.

The continuous formulation of intrinsic dimension has a number of potential applications. For example, in optic flow analysis it can be used to distinguish between normal flow (a i1D signal patches) and potentially correct flow (at i2D image patches). The continuous formulation could allow for an appropriate weighting of flow vectors for global optic flow interpretation. Another example is the accumulation of ambiguous information over time (see, e.g., [23]). The continuous formulation would allow for the postponing of a final decision about edge-ness or junction-ness to a rather later stage of processing that can make use of a number of time frames.

An extension of this work, in which the triangle formulation is extended to a cone representation allowing for a probabilistic of image patches is described in [12].

Acknowledgment: We would like to thank Florentin Wörgötter for fruitful discussions. This work has been funded by the European Project ECOVISION.

References

- [1] J. Aloimonos and D. Shulman. *Integration of Visual Modules — An extension of the Marr Paradigm*. Academic Press, London, 1989.
- [2] H. Bärman, G. H. Granlund, and H. Knutsson. Tensor Field Filtering and Curvature Estimation. In *Proceedings of the SSAB Symposium on Image Analysis*, pages 175–178, Linköping, Sweden, March 1990. SSAB. Report LiTH-ISY-1-1088, Linköping University, Sweden, 1990.

- [3] J. Bigün and G. H. Granlund. Optimal orientation detection of linear symmetry. In *Proceedings of the IEEE First International Conference on Computer Vision*, pages 433–438, London, Great Britain, June 1987. Report LiTH-ISY-I-0828, Computer Vision Laboratory, Linköping University, Sweden, 1986.
- [4] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, New York, 1995.
- [5] R. N. Bracewell. *The Fourier transform and its applications*. McGraw Hill, 1986.
- [6] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6), 1986.
- [7] H.S.M. Coxeter. *Introduction to Geometry (2nd ed.)*. Wiley & Sons, 1969.
- [8] S. C. Dakin. Local orientation variance as a quantifier of structure in texture. *Spatial Vision*, 12:1–30, 1999.
- [9] ECOVISION: Artificial visual systems based on early-cognitive processing, EU-Project. <http://www.pspc.dibe.unige.it/ecovision/project.html>, 2002–2004.
- [10] G. Farneböck. Fast and accurate motion estimation using orientation tensors and parametric motion models. *Proc. ICPR*, 2000.
- [11] M. Felsberg. *Low-Level Image Processing with the Structure Multivector*. PhD thesis, Institute of Computer Science and Applied Mathematics, Christian-Albrechts-University of Kiel, 2002.
- [12] M. Felsberg and N. Krüger. A probabilistic definition of intrinsic dimensionality for images. *Pattern Recognition, 24th DAGM Symposium*, 2003.
- [13] M. Felsberg and G. Sommer. Image features based on a new approach to 2D rotation invariant quadrature filters. *Proc. of ECCV*, pages 369–383, 2000.
- [14] M. Felsberg and G. Sommer. The monogenic signal. *IEEE Transactions on Signal Processing*, 49(12):3136–3144, December 2001.
- [15] W. Förstner. *Statistische Verfahren für die automatische Bildanalyse und ihre Bewertung bei der Objekterkennung und -vermessung*. Number 370 in C. Verlag der Bayerischen Akademie der Wissenschaften, 1991.
- [16] W. Förstner and E. Gülch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *ISPRS Intercommission Workshop, Interlaken*, pages 149–155, June 1987.
- [17] G. H. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, Dordrecht, 1995.
- [18] C. G. Harris and M. Stephens. A combined corner and edge detector. In *4th Alvey Vision Conference*, pages 147–151, 1988.
- [19] B. Jähne. *Digital Image Processing – Concepts, Algorithms, and Scientific Applications*. Springer, 1997.
- [20] J.J Koenderink. What is a feature? *J. Intell. Syst.*, 3(1):49–82, 1993.
- [21] N. Krüger. Collinearity and parallelism are statistically significant second order relations of complex cell responses. *Proceedings of I&ANN 98*, 1998.
- [22] N. Krüger. Learning object representations using a priori constraints within ORASSYLL. *Neural Computation*, 13(2):389–410, 2001.
- [23] N. Krüger, M. Ackermann, and G. Sommer. Accumulation of object representations utilizing interaction of robot action and perception. *Knowledge Based Systems*, 15:111–118, 2002.
- [24] N. Krüger, M. Felsberg, and F. Wörgötter. Processing multi-modal primitives from image sequences. *submitted to BMVC 2003*, 2003.
- [25] P. Parent and S.W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):823–839, 1989.
- [26] E. Ribeiro and E.R. Hancock. Shape from periodic texture using the eigenvectors of local affine distortion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(12):1459–1465, 2001.
- [27] K. Rohr. Recognizing corners by fitting parametric models. *International Journal of Computer Vision*, 9(3):213–230, 1992.
- [28] G. V. Trunk. Representation and analysis of signals: statistical estimation of intrinsic dimensionality and parameter identification. *General System*, 13:49–76, 1968.
- [29] C. Zetsche and E. Barth. Fundamental limits of linear filters in the visual processing of two dimensional signals. *Vision Research*, 30, 1990.