

Optic Flow Statistics and Intrinsic Dimensionality

Sinan Kalkan¹, Dirk Calow², Michael Felsberg³, Florentin Wörgötter¹, Markus Lappe², and Norbert Krüger⁴

¹ Psychology, University of Stirling, Scotland
sinan,worgott@cn.stir.ac.uk

² Psychology, University of Münster, Germany
calow,mlappe@psy.uni-muenster.de

³ Electrical Eng., Linköping University, Sweden
mfe@isy.liu.se

⁴ Computer Sci., Aalborg University Esbjerg, Denmark
nk@cs.aue.auc.dk

ABSTRACT

Different kinds of visual sub-structures (such as homogeneous, edge-like and junction-like patches) can be distinguished by the intrinsic dimensionality of the local signals. The concept of intrinsic dimensionality has been mostly exercised using discrete formulations. A recent work [KF03,FK03] introduced a continuous definition and showed that the inherent structure of the intrinsic dimensionality has essentially the form of a triangle. The current study work analyzes the distribution of signals according to the continuous interpretation of intrinsic dimensionality and the relation to orientation and optic flow features of image patches. Among other things, we give a quantitative interpretation of the distribution of signals according to their intrinsic dimensionality that reveals specific patterns associated to established sub-structures in computer vision. Furthermore, we link quantitative and qualitative properties of the distribution of optic-flow error estimates to these patterns.

1 INTRODUCTION

Natural images are dominated by specific local sub-structures, such as homogeneous patches, edges, corners, or textures. Sub-domains of Computer Vision have extracted and analyzed such sub-structures in edge detection (see, e.g., [Can86]), junction classification (see, e.g., [Roh92]) and texture interpretation (see, e.g., [RH01]).

The intrinsic dimension (see, e.g., [ZB90,Fel02]) has proven to be a suitable descriptor that distinguishes such sub-structures. 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). The association of intrinsic dimension to a local image structure has been done mostly by a discrete classification [ZB90,Fel02,J97]. A *continuous* definition of intrinsic dimensionality has been recently given in [KF03,FK03]. There, it has also been shown that *the topological structure of intrinsic dimension essentially has the form of a triangle* which is spanned by two axis corresponding to origin and line variance. In this paper, we will use this continuous definition to investigate the structure of the distribution of signals in natural images

according to their intrinsic dimensionality. More specifically, we will show that;

- D0 i0D signals split into two clusters. One peak corresponding to saturated or dark patches and a Gaussian-shaped cluster corresponding to image noise at homogeneous but unsaturated/non-black image patches.
- D1 For i1D signals, there exists a concentration of signals in a stripe-shaped cluster corresponding to high origin variance (high amplitude) and low line variance. This reflects the importance of an orientation criterion that is based on local amplitude and orientation information (see, e.g., [PIK90]).
- D2 In contrast to the i0D and i1D cases, there exists no such thing like a corner cluster for i2D signals in the distribution of local signals which indicates that it is rather difficult to formulate a local criterion to detect corners in natural images.

Thus, the continuous formulation of intrinsic dimension allows for a more precise characterization of established sub-structures in terms of their statistical manifestation in natural images. As a consequence, properties and inherent problems of classical computer vision algorithms can be reflected, and the limits of local signal processing can be made explicit.

The property of optic flow estimation at homogeneous image patches, edges, and corners has been discussed extensively (see, e.g., [BFB94,MB00]). In general, it is acknowledged that;

- A0 Optic flow estimates at homogeneous image patches is unreliable due to the fact that the lack of structure makes it impossible to find correspondences in consecutive frames.
- A1 Optic flow at edge-like structures faces the aperture problem such that only the normal flow can be computed for these structures.
- A2 Only optic flow estimation at i2D structures can lead to true optic flow estimates.

In this paper, we investigate these claims more closely. We will show that the continuous formulation of intrinsic dimensionality allows for a more quantitative investigation and characterization of the quality of optic flow estimation depending on local signal structures. More specifically, we will show that;

- Q0 There exist significantly more horizontal and vertical structures in natural images. However, the strength of the dominance of these structures depends crucially on the intrinsic dimension. Furthermore, we show that the distribution of orientations is directly reflected in the distribution of the estimates of optic-flow directions.
- Q1 Optic flow can be estimated reliably when looking at the normal flow in the stripe-shaped cluster in the i1D signal domain. This suggests that edges are a strong source of reliable information when the aperture problem is taken into account.
- Q2 The quality of optic flow estimation is higher for i2D signals. However, in analogy to the lack of a cluster for i2D signals, there exists a continuous signal domain (covering also sub-areas of i0D and i1D signals) for which a higher quality in the optic flow estimation can be achieved. The increase of the quality, on the other hand, is only slight which suggests that the role of i2D structures for motion estimation might be not as important as suggested by some authors (see, e.g., [MB00]). Another possibility is that the specific optic flow algorithm that we have chosen gives sub-optimal results at i2D signals.

These results support the above-mentioned statements (A0)-(A2) about optic flow estimation. However, by making use of a continuous understanding of intrinsic dimensionality, these statements can be made quantitatively more specific in terms of (1) characterization of sub-areas for which they hold and (2) the strength of these statements. Our analysis suggests a strong relationship between the distribution of the signals in the continuous intrinsic dimensionality space and properties of optic flow estimation.

The outline of the paper is as follows: In section 2, we introduce the concept of intrinsic dimensionality in its continuous definition as first formulated in [KF03,FK03] and we shortly describe our multi-modal image representation in which orientation and optic flow are coded. In section 3, we investigate the distribution of local image patches according to their intrinsic dimensionality and discuss its consequences for image processing. In section 4, we look at the statistics of orientation in natural images while in section 5, we look at the statistics of optic flow and the error of optic flow estimation.

2 INTRINSIC DIMENSIONALITY AND MULTI-MODAL PRIMITIVES

Intrinsic Dimensionality: In image processing, the intrinsic dimensionality has been used to characterize local image patches according to their homogeneousness, edge-ness or junction-ness. The term intrinsic dimensionality is itself much more general.

Accordingly, when looking at the spectral representation of a local image patch (see figure 1a,b), we see that the energy of an intrinsically zero-dimensional signal is concentrated in the origin (figure 1b-top), the energy of an intrinsically one-dimensional signal is concentrated along a line (figure 1b-middle) while the energy of an intrinsically two-dimensional signal varies in more than one dimension (figure 1c-bottom).




Fig. 1. Illustration of intrinsic dimensionality. (a) Three image patches for three different intrinsic dimensions. (b) The local spectra of the patches in (a), from top to bottom: i0D, i1D, i2D. (c) The topology of intrinsic dimensionality. (d) The primitives shown on an image used in our analysis. Red, green and blue colors show i0D, i1D and i2D primitives, respectively.

In image processing, the intrinsic dimensionality was introduced by [ZB90] and was used to formalise a *discrete distinction* between edge-like and junction-like structures. This corresponds to a classical interpretation of local image signals in computer vision. A large variety of edge and corner extraction algorithms have been developed over the last 20 years, and their role in artificial as well as biological systems has been discussed extensively [KD82]. The three kinds of signals have quite specific characteristics and problems that have been addressed in different contexts.

Recently, it has been shown [KF03,FK03] that the topological structure of the intrinsic dimensionality must be understood as a triangle that is spanned by two measures: origin variance and line variance. The origin variance describes the deviation of the energy from a concentration at the origin while the line variance describes the deviation from a line structure (see figure 1b and 1c).

The triangular topological structure of the intrinsic dimensionality is counter-intuitive, in the first place, since it realizes a two-dimensional topology in contrast to a linear one-dimensional structure that is expressed in the discrete counting 0, 1 and 2. More importantly, as shown in [KF03,FK03], this triangular interpretation allows for a *continuous formulation* of intrinsic dimensionality in terms of 3 confidences

assigned to each discrete case. This is achieved by first computing two measurements of origin and line variance. These two measurements define a point in the triangle (see figure 1c). The bary-centric coordinates (see, e.g., [Cox69]) of this point of the triangle directly lead to a definition of three confidences that add up to one. These three confidences reflect the volume of the 3 areas constituted by the points (see figure 1c).

Multi-modal Primitives: In our research over the last 5 years [KLW04], we have designed a novel image representation in terms of multi-modal Primitives (see figure 2). They represent a condensed representation of local image structure in terms of visual attributes that are also found in the first stages of human visual processing (see figure 2). These Primitives are a functional abstraction of hyper-columnar structures found in the first cortical area of the human visual system [HW69]. These Primitives carry, beside the 3 confidences for the three intrinsic dimensions of a local signal also other attributes such as orientation, contrast transition, colour and optic flow.




Fig. 2. Left: Image sequence and frame. Middle: Schematic representation of the multi-modal Primitives. Right: Extracted Primitives at positions with high amplitude.

Orientation: The local orientation associated to a Primitive is computed using a recently developed filter approach, called the monogenic signal [FS01]. The orientation is computed by interpolating across the orientation information of the whole image patch to achieve a more reliable estimate.

Optic Flow: There exist a large variety of algorithms that compute the local displacement in image sequences. [BFB94] have divided them into 4 classes: differential techniques, region-based matching, energy based methods and phase-based techniques. After some comparison [JÖ2], we decided to use the well-known optic flow technique [NE86]. This algorithm is a differential technique in which, however, in addition to the standard gradient constraint equation, an anisotropic smoothing term is supposed to lead to better flow estimation at edges (for details see [NE86,BFB94,AWS00]).

3 STATISTICS OF INTRINSIC DIMENSIONALITY

We use a set of 7 natural sequences with 10 images each (see figure 3). We compute 831727 Primitives in total. For



Fig. 3. Some of the image sequences used in our analysis. The first 3 images are from one of the sequences (the starting image, the middle image and the last image). Other figures are the images from other sequences.

the processing of each Primitive, we compute a measure for the origin and line variance (for details see [KF03]). This corresponds to one point in the triangle of figure 1c. Hence, taking all Primitives into account, we can display the distribution of these points in this triangular structure. This distribution is shown in figure 4a. As there exist large differences in the histogram, the logarithm is shown.

The distribution shows two main clusters. The peak close to the origin corresponds to low origin variance. It is visible that most of the signals that have low origin variance have high line variance. These correspond to nearly homogeneous image patches. Since the orientation is almost random for such homogeneous image patches, it causes high line variance. There is also a small peak existent at position $(0,0)$ that corresponds to saturated/black image patches. The other cluster is for high origin variance signals with low line variance, corresponding to edge-like structures. The form of this cluster is a small horizontal stripe rather than a peak. There is a smooth decrease while approaching to the i2D area of the triangle. That means that there exists no cluster for corner-like structures like the ones for homogeneous image patches or edges. Along the origin variance axis, a continuous gap is observed. This gap suggests that there are no signals with zero variance. This is due to the fact that there is at least noise included in the image which causes some line variance.

We also see from the figure that there are far more i0D signals than i1D or i2D signals. So, we have much more homogeneous structures than edge-like or junction-like structures. Besides, it is clear that there are more i1D structures than i2D structures in natural images. The percentages of i0D, i1D and i2D structures are %86, %11 and %3, respectively. For computing these percentages, the intrinsic dimensionality of a Primitive is determined by taking the intrinsic dimension which has the highest confidence. Figure 1d shows a natural scene on which Primitives are shown. Due to space constraints, only a small portion of an image in a sequence is shown. The percentages of i0D, i1D and i2D structures are clearly reflected in the figure.

The results described above are taken from Primitives for which the position has been regularly sampled on a hexagonal grid. However, the position of an edge or a corner must be defined according to the internal structure of the




Fig. 4. Logarithmic plot of the distribution of intrinsic dimensionality. **(a)** The distribution when the positions of the Primitives are not modified according to iD. **(b)** The distribution when the positions are modified according to iD (See text for details of this modification).

signal. For example, we want an edge-like Primitive to be placed directly on the edge; or, for a corner, in which a certain number of lines intersect, we want to have the Primitive placed on the intersection. We can achieve this positioning by making use of the local amplitude information in the image depending on the intrinsic dimensionality (for details, see [KFW04]). Note also that features such as phase and the optic flow depends on this positioning. When we determine the position of the Primitives for edges and corner-like structures accordingly, we get the distribution shown in figure 4b. It is qualitatively similar to the distribution achieved with regular sampling. However, since the position is determined depending on the local amplitude, there is a shift towards positions with higher amplitude that constitute the gaps at the border between i0D and the i1D and i2D signals. In the later stages of our analysis, we adopted this positioning.

4 DISTRIBUTION OF ORIENTATION DEPENDENT ON INTRINSIC DIMENSIONALITY

The distribution of the orientation on the image sequences and the quantitative differences depending on the intrinsic

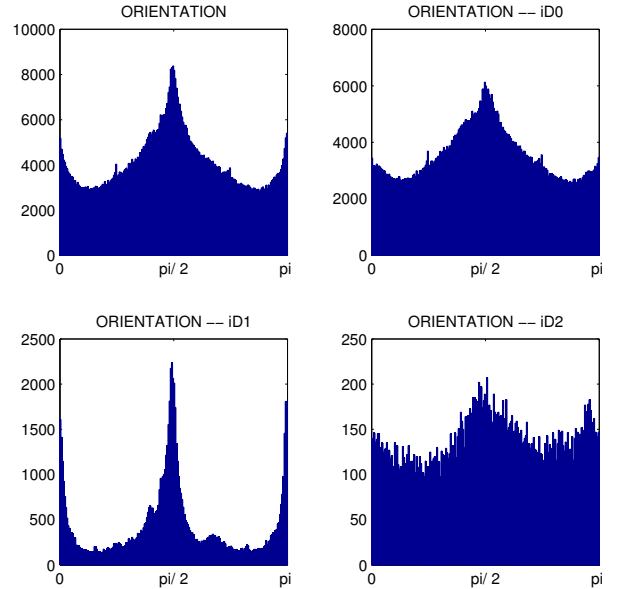


Fig. 5. Orientation distribution depending on iD. The first image shows the total distribution.

dimensionality of the Primitives are shown in figure 5. Here, we define a signal to be i_nD (where n is 0, 1 or 2) if the associated confidence is the highest. We see that orientation of $i1D$ signals shows strong peaks at horizontal and vertical structures (i.e., for the values 0, $\pi/2$ and π). This has been already known (see, e.g., [Krü98]). However, the statistics in figure 5 for the $i0D$ and $i2D$ cases show the quantitative dependency of this statement for signals with different intrinsic dimension: These peaks are much weaker for the $i0D$ and $i2D$ case. Indeed, neither for a completely homogeneous image patch nor for a corner, the concept of orientation is defined at all. However, the continuous formulation of intrinsic dimensionality prevails that as dominance of horizontal and vertical orientations can be found also for $i0D$ and $i2D$ signals. This means that orientation is a meaningful concept also for some non- $i1D$ signals. This also stresses the advantages of a continuous understanding of intrinsic dimensionality.

5 OPTIC FLOW AND INTRINSIC DIMENSIONALITY

The distribution of magnitude and direction of the optic flow vectors is shown in figure 6. The quantitative errors in calculation of optic flow are shown in 8a-f using three different measurements.

The distribution of direction varies significantly with the intrinsic dimensionality. The distribution of the direction of optic flow vectors of $i1D$ signals directly reflects the statistics of orientation. Since only the normal flow can be computed for pure $i1D$ signals, the dominance of horizontal and vertical orientations (see section 4) leads to peaks at horizontal and vertical flow. The statistics of the true flow can be expected to be nearly homogeneous since in the sequences, a translational forward motion is dominant that leads to an isotropic flow field (see, e.g., [LBv99]). The fact that basically there exists a direct quantitative equivalence of the statistics of line structures, and the statistics of optic flow

directions reflects the seriousness of the aperture problem. In contrast, the distribution of direction of optic flow vectors of i0D and i2D signals is nearly homogeneous.

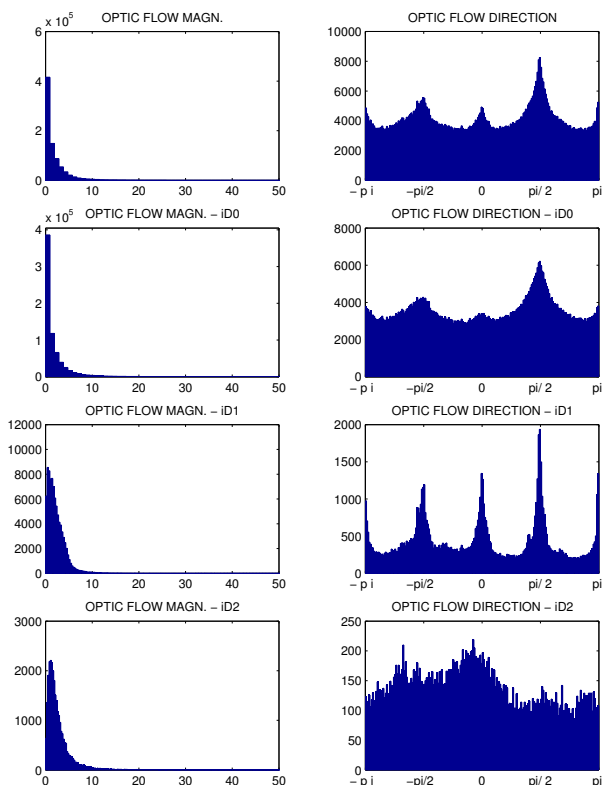


Fig. 6. Optic flow distributions depending on ID. The left two images show total distributions.

We now analyze the errors of the optic flow estimation depending on the intrinsic dimension. For this, we need to compare the computed flow with a ground truth. We used for our investigations the Brown Range Image Database (brid), a database of 197 range images collected by Ann Lee, Jing-gang Huang and David Mumford at Brown University (see also [HLM00]). The range images are recorded with a laser range-finder. Each image contains 44×1440 measurements with an angular separation of 0.18 degree. The field of view is 80 degree vertically and 259 degree horizontally. The distance of each point is calculated from the time of flight of the laser beam, where the operational range of the sensor is 2–200m. The laser wavelength of the laser beam is $0.9\mu\text{m}$ in the near infrared region. Thus, the data of each point consist of 4 values, the distance, the horizontal angle and the vertical angle in spherical coordinates and a value for the reflected intensity of the laser beam (see figure 7). The knowledge about the 3D data structure allows for a simulation of a moving camera in a scene and can therefore be used to estimate the correct flow for nearly all pixel positions of a frame of an image sequence.

The error of optic flow calculation is displayed in a histogram over the iD triangle (see figure 8). When calculating the error, three different error functions were used: Let \mathbf{u} be the estimated and \mathbf{v} be the correct optic flow vector; then, the angular error is computed by $e(\mathbf{u}, \mathbf{v}) = \text{acos}(\frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}| |\mathbf{v}|})$ (figure 8b,e); the magnitudal error using $e(\mathbf{u}, \mathbf{v}) = \text{abs}(|\mathbf{u}| - |\mathbf{v}|)$




Fig. 7. Sequences that we have used in the statistics of error estimates. The first three images are from the same sequences (the starting image, the middle image and the last image). Other images are from other sequences.

(figure 8c,f); and a mixture of both (see also [BFB94]) using $e(\mathbf{u}, \mathbf{v}) = \text{acos}(\frac{\mathbf{u} \cdot \mathbf{v} + 1}{(\mathbf{u} \cdot \mathbf{u} + 1)(\mathbf{v} \cdot \mathbf{v} + 1)})$ (figure 8a,d).

Since for purely i1D signals, the aperture problem allows for the computation of the normal flow only, we have also computed errors with the projection of the ground truth over the normal vectors of the Primitives. For this, we first compute the normal vector of the Primitive using its orientation information; the ground truth is projected over this vector, and the error is computed between the optic flow of the Primitive and this projected ground truth.

The combined error computed using the original ground truth (see figure 8a) is high for signals close to the i0D corner of the triangle as well as on the horizontal stripe from the i0D and i1D corner. In the other parts, there is a smooth surface that slightly decreases towards the i2D corner. This is in accordance with the notion that optic flow estimation at corner-like structures is more reliable than for edges and homogeneous image patches (A2). However, in figure 8a, it becomes obvious that the area where more reliable flow vectors can be computed is very broad and covers also traditional i0D and i1D signals. Furthermore, the decrease of error is rather slight which points to the fact that the quality of flow computation is limited in these areas, as well.

The current study makes use of a specific optic flow algorithm [NE86]. To test whether the slight decrease of the error for i2D signals is a property of Nagel algorithm or a general feature of other optic flow algorithms as well, we used the Lucas Kanade optic flow algorithm [LK81] on the same sequences used in our quality analysis. The results are shown in figure 9a-f. We see from the figure 9a that the area with low optic flow error extends to areas of i1D and i0D signals with small changes. This suggests that this only slight decrease is a general property of optic flow algorithms rather than a specific property of the Nagel algorithm.

For the combined error computed using the normal ground truth (see figure 8d), a different picture occurs. The error is very low for a horizontal stripe from the middle point between the i0D and i1D corners to the i1D corner. When compared to figure 8a, this figure reflects the strength of the aperture problem. On the other hand, it also shows the quality of optic flow estimation when the aperture problem is taken into account. The information for such signals can be of great importance, since for example, constraints for global motion estimation can be defined on line correspondences (see, e.g., [Ros03,KW04]), i.e., correspondences that only require normal flow. In figure 8, the angular and magnitude

The Nagel Algorithm

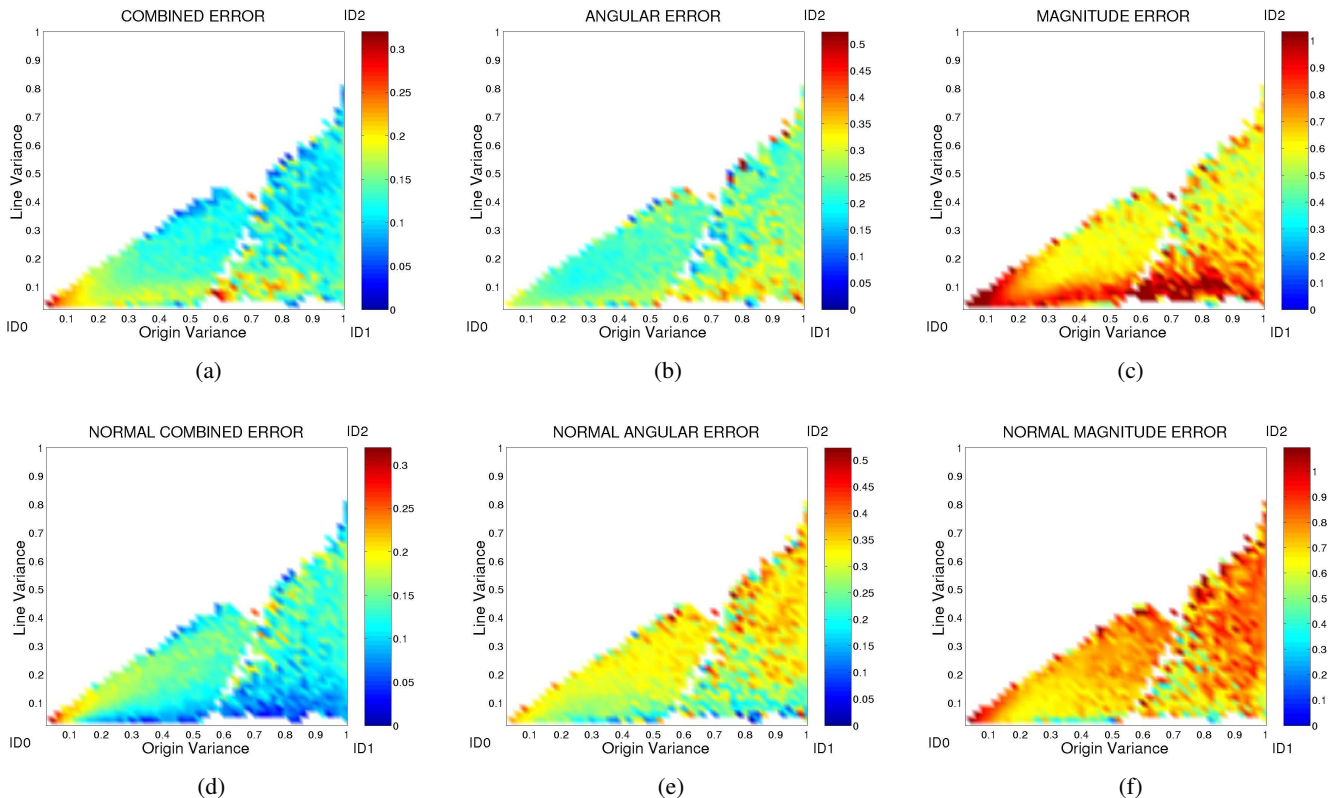


Fig. 8. Qualities of the Nagel optic flow algorithm depending on iD . Color bars show the error values for corresponding colors of corresponding graphs. **(a)-(c)** Errors with the optic flow ground truth. **(d)-(f)** Errors with the projection of the ground truth over the normal vectors of the Primitives. **(a,d)** Errors computed using a mixture of magnitudes and angles. **(b,e)** Errors computed using angles. **(c,f)** Errors computed using magnitudes.

errors are also displayed that reflect a similar behaviour, especially visible for the angular error.

6 DISCUSSIONS

Using statistics over natural image sequences, we analyzed the distribution of local signal patches according to their intrinsic dimensionality. Furthermore, we investigated how the quality of optic flow estimates and orientation of image patches depend on intrinsic dimensionality.

A continuous understanding of the concept of intrinsic dimension [KF03,FK03] allows for a more precise characterization of established sub-structures in terms of their statistical manifestation in natural images. The continuous formulation of intrinsic dimensionality also allows for a more quantitative investigation and characterization of the quality of optic flow estimation depending on local signal structures. We could justify and more precisely quantify generally acknowledged ideas about such estimates. Moreover, we could also point to rather surprising results for the $i2D$ case.

In general, we could show that such a continuous understanding of intrinsic dimension reflects the statistical properties in natural images: The distribution of signals in the iD -triangle is rather smooth with some specific sub-clusters. Besides, the quality of the flow estimates varies continuously with the intrinsic dimension.

The statistics about the $i1D$ case made the role of the aperture problem explicit. The most reliable flow estimates

can be achieved for a certain subset of $i1D$ signals. However, then the aperture problem needs to be accounted for.

REFERENCES

- [AWS00] L. Alvarez, J. Weickert, and J. Sanchez. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 39:41–56, 2000.
- [BFB94] J.L. Barron, D.J. Fleet, and S.S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994.
- [Bis95] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, New York, 1995.
- [Can86] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6), 1986.
- [Cox69] H.S.M. Coxeter. *Introduction to Geometry (2nd ed.)*. Wiley & Sons, 1969.
- [Fel02] 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.
- [FK03] M. Felsberg and N. Krüger. A probabilistic definition of intrinsic dimensionality for images. *Pattern Recognition, 24th DAGM Symposium*, 2003.
- [FS01] M. Felsberg and G. Sommer. The monogenic signal. *IEEE Transactions on Signal Processing*, 49(12):3136–3144, December 2001.
- [HLM00] J. Huang, A.B. Lee, and D. Mumford. Statistics of range images. *CVPR*, pages 1324–1331, 2000.

The Lucas-Kanade Algorithm




Fig. 9. Qualities of the Lucas-Kanade optic flow algorithm depending on iD. Color bars show the error values for corresponding colors of corresponding graphs. **(a)-(c)** Errors with the optic flow ground truth. **(d)-(f)** Errors with the projection of the ground truth over the normal vectors of the Primitives. **(a,d)** Errors computed using a mixture of magnitudes and angles. **(b,e)** Errors computed using angles. **(c,f)** Errors computed using magnitudes.

- [HW69] D.H. Hubel and T.N. Wiesel. Anatomical demonstration of columns in the monkey striate cortex. *Nature*, 221:747–750, 1969.
- [J97] B. Jähne. *Digital Image Processing – Concepts, Algorithms, and Scientific Applications*. Springer, 1997.
- [J02] Thomas Jäger. Interaktion verschiedener visueller Modalitäten zur stabilen Extraktion von objektrepräsentationen. *Diploma thesis (University of Kiel)*, 2002.
- [KD82] J.J. Koenderink and A.J. Dorn. The shape of smooth objects and the way contours end. *Perception*, 11:129–173, 1982.
- [KF03] N. Krüger and M. Felsberg. A continuous formulation of intrinsic dimension. *Proceedings of the British Machine Vision Conference*, 2003.
- [KFW04] N. Krüger, M. Felsberg, and F. Wörgötter. Processing multi-modal primitives from image sequences. *Fourth International ICSC Symposium on ENGINEERING OF INTELLIGENT SYSTEMS*, 2004.
- [KLW04] N. Krüger, M. Lappe, and F. Wörgötter. Biologically motivated multi-modal processing of visual primitives. *The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour*, 1(5), 2004.
- [Krü98] N. Krüger. Collinearity and parallelism are statistically significant second order relations of complex cell responses. *Proceedings of I&ANN 98*, 1998.
- [KW04] N. Krüger and F. Wörgötter. Statistical and deterministic regularities: Utilisation of motion and grouping in biological and artificial visual systems. *Advances in Imaging and Electron Physics*, 131, 2004.
- [LBv99] M. Lappe, F. Bremmer, and A. V. van den Berg. Perception of self-motion from visual flow. *Trends in Cognitive Sciences*, 3:329–336, 1999.
- [LK81] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *Proc. DARPA Image Understanding Workshop*, pages 121–130, 1981.
- [MB00] C. Mota and E. Barth. On the uniqueness of curvature features. *Proc. in Artificial Intelligence*, 9:175–178, 2000.
- [NE86] H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
- [PIK90] J. Princen, J. Illingworth, and J. Kittler. An optimizing line finder using a Hough transform algorithm. *Computer Vision, Graphics, and Image Processing*, 52:57–77, 1990.
- [RH01] 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.
- [Roh92] K. Rohr. Recognizing corners by fitting parametric models. *International Journal of Computer Vision*, 9(3):213–230, 1992.
- [Ros03] B. Rosenhahn. *Pose Estimation Revisited (PhD Thesis)*. Institut für Informatik und praktische Mathematik, Christian-Albrechts-Universität Kiel, 2003.
- [ZB90] C. Zetsche and E. Barth. Fundamental limits of linear filters in the visual processing of two dimensional signals. *Vision Research*, 30, 1990.