Statistical Analysis of Local 3D Structure in 2D Images

Sinan KALKAN
Bernstein Centre for Computational Neuroscience,
University of Göttingen, Germany
sinan@chaos.gwdg.de

Florentin Wörgötter
Bernstein Centre for Computational Neuroscience,
University of Göttingen, Germany
worgott@chaos.gwdg.de

Norbert Krüger
Cognitive Vision Group,
Aalborg University Copenhagen, Denmark
nk@media.aau.dk

Abstract

*For the analysis of images, a deeper understanding of their intrinsic structure is required. This has been obtained for 2D images by means of statistical analysis [15, 18]. Here, we analyze the relation between local image structures (*i.e.*, homogeneous, edge-like, corner-like or texture-like structures) and the underlying local 3D structure, represented in terms of continuous surfaces and different kinds of 3D discontinuities, using 3D range data with the true color information. We find that homogeneous image patches correspond to continuous surfaces, and discontinuities are mainly formed by edge-like or corner-like structures. The results are discussed with regard to existing and potential computer vision applications and the assumptions made by these applications.*

1. Introduction

With the notion that the human visual system is adapted to the statistics of the environment [2, 13, 15, 18, 22, 21] and its successful applications to grouping, object recognition and stereo [3, 4, 20, 29] the analysis, and the usage of natural image statistics has become an important focus of vision research. Moreover, with the advances in technology, it has been also possible to analyze the underlying 3D world using 3D range scanners [10, 11, 19, 27].

In this paper, we analyze the relation between local image structures (*i.e.*, homogeneous, edge-like, corner-like or texture-like structures) and the underlying local 3D structure using 3D range data with the true color information.

There have been only a few studies that have analyzed the 3D world from range data [10, 11, 19, 27]. In [27], the distribution of roughness, size, distance, 3D orientation, curvature and independent components of surfaces was analyzed. Their major conclusions were: (1) local 3D patches tend to be saddle-like, and (2) natural scene geometry is quite regular and less complex than luminance images. In [11], the distribution of 3D points was analyzed using co-occurrence statistics and 2D and 3D joint distributions of Haar filter reactions. They showed that range images are much simpler to analyze than optical images and that a 3D scene is composed of piecewise smooth regions. In [19], the correlation between light intensities of the image data and the corresponding range data as well as surface convexity were investigated. They could justify the event that brighter objects are closer to the viewer, which is used by shape from shading algorithms in estimating depth. In [9, 10], range image statistics were analyzed for explanation of several visual illusions.

Our analysis differs from these works. For 2D local image patches, existing studies have only considered light intensity. As for 3D local patches, the most complex considered representation have been the curvature of the local 3D patch. In this work, however, we create a higher-order representation of the 2D local image patches and the 3D local patches; we measure 2D local image patches using homogeneous, edge-like, corner-like or texture-like structures, and 3D local patches using continuous surfaces and different kinds of 3D discontinuities. By this, we relate established local image structures to their underlying 3D structures.

By creating 2D and 3D representations of the local structure, we compute the conditional probability $P(\text{3D Structure} \mid \text{2D Structure})$. Using this probability, we quantify some assumptions made by the studies that reconstruct the 3D world from dense range data. For example, we could show that the depth distribution varies significantly for different visual features, and we could quantify already established inter-dependencies such as 'no new is

good news' [6]. This work also supports the understanding of how intrinsic properties 2D–3D relations can be used for the reconstruction of depth, for example, by using statistical priors in the formalisation of depth cues.

The paper is organized as follows: In section 2, we define the types of local image structures and local 3D structures that we extract for our analysis. In section 3, we introduce a continuous classifier for local 2D structures. In section 4, we outline our methods for measuring the 3D structure of a 3D point. We present and discuss our results in section 5. Finally, we conclude the paper in section 6.

2. Local 2D and 3D Structures

We distinguish between the following local 2D structures:

- Homogeneous image patches: Homogeneous patches are signals of uniform intensities.
- Edge–like structures: Edges are low-level structures which constitute the boundaries between homogeneous or texture-like signals (see, *e.g.*, [14, 17] for their importance in vision).
- Corners: Corners are signals where two or more edge-like structures with significantly different orientations intersect (see, *e.g.*, [7, 23, 24] for their importance in vision).
- Texture: Although there is not a widely-agreed definition, textures are often defined as signals which consist of repetitive, random or directional structures (for their analysis, extraction and importance in vision, see *e.g.*, [26]).

Locally, it is hard to distinguish between these structures, and there are structures that carry mixed properties of the 'ideal' cases. The classification of the features outlined above is discrete. However, a discrete classification may cause problems as the inherent properties of "mixed" structures are lost in the discretization process. Instead, in this paper, we make use of a recently developed continuous scheme which is based on the concept of intrinsic dimensionality [5, 16]. In this concept, local image structures are organized continuously in a triangle. This approach is briefly described in section 3. Here, we show that the different classes of local image structures map to different distinguishable areas in the domain of the intrinsic dimensionality triangle (see figure 2) which is the first contribution of this paper.

To our knowledge, there does not exist a systematic and agreed classification of 3D local structures like there is for 2D local image structures (*i.e.*, homogeneous patches, edges, corners and textures). Intuitively, the 3D world consists of continuous surface patches and different kinds of 3D discontinuities. In the imaging process (through the lenses of camera or a retina), 2D local image structures are formed by these 3D structures together with the illumination and reflectivity of the environment.

With this intuition, any 3D scene can be decomposed geometrically into surfaces and 3D discontinuities. In this context, the local 3D structure of a point can be a:

- Surface Continuity: The underlying 3D structure can be described by one surface whose normal does not change or changes smoothly.
- Regular Gap discontinuity: The underlying 3D structure can be described by a small set of surfaces with a significant depth difference. The 2D and 3D views of an example gap discontinuity are shown in figure 1(a).
- Irregular Gap discontinuity: The underlying 3D structure shows high depth variation and can not be described by two or three surfaces. An example of an irregular gap discontinuity is shown in figure 1(b).
- Orientation Discontinuity: The underlying 3D structure can be described by two surfaces with significantly different 3D orientations that meet at the point whose 3D structure is being questioned. In this type of discontinuity, no gap but a change in 3D orientation between the meeting surfaces occurs. An example for this type of discontinuity is shown in figure 1(c).

3. Intrinsic Dimensionality

In image processing, intrinsic dimensionality was introduced by Zetsche and Barth[28] to distinguish between different local image structures. The idea is to assign intrinsically zero dimensionality (i0D), intrinsically one dimensionality (i1D) and intrinsically two dimensionality (i2D) to homogeneous patches, edges and corner-like structures, respectively. The concept of intrinsic dimensionality has been mostly applied in a discrete way which has been extended in [5, 16] to classify the local image structures continuously instead of giving them discrete labels.

In [5, 16], it has been also shown that the topological structure of the intrinsic dimensionality can be understood as a triangle whose corners correspond to the 'ideal' cases of 2D structures (*i.e.*, homogeneous patches, edges and corners). The inner of the triangle spans signals that carry aspects of the three 'ideal' cases, and the distance from the specific corners indicates the similarity (or dissimilarity) to the 'ideal' i0D, i1D and i2D signals. The horizontal and the vertical axes denote the contrast and the orientation variance, respectively. Contrast measures non-homogeneity whereas orientation variance measures the variation of orientation in a local patch describing the local image structure. An 'ideal' homogeneous image patch is expected to have zero contrast and zero orientation variance whereas an 'ideal' edge should have high contrast and zero orientation variance. An 'ideal' corner is supposed to have high contrast and high orientation variance.

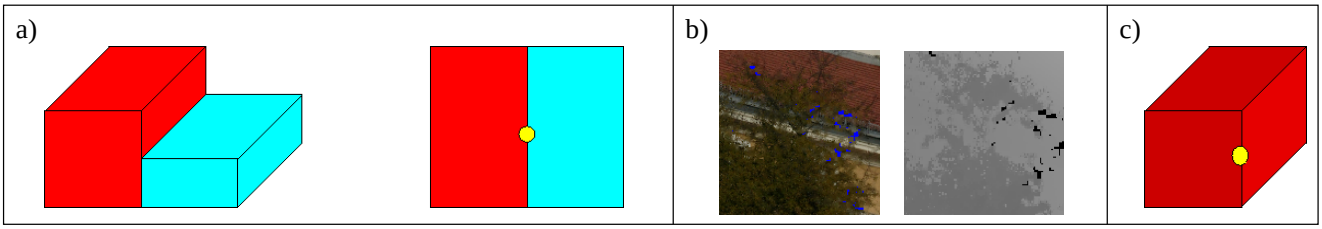

Figure 1. Examples for types of 3D discontinuities. Points of interest are marked with yellow circles. (a) 2D and 3D views of a gap discontinuity, (b) image (on the left) and range data (on the right) of an irregular gap discontinuity and (c) orientation discontinuity.

Figure 2 shows how the triangle of intrinsic dimensionality looks like and how a set of example local image structures map on to it. In figure 2, we see that different visual structures map to different areas in the triangle. A detailed analysis of how 2D structures are distributed over the intrinsic dimensionality triangle and how some visual information depends on this distribution can be found in [12]. Different from [12], in this paper, regarding this distribution, we show that textures also map to a different area of their own. The fact that different local image structures have their own distinguishable areas in the triangle provides us with a continuous classifier that distinguishes between homogeneous, edge-like, texture-like and corner-like structures.

4. Methods

In this section, we define our measures for the three kinds of discontinuities that we described in section 1; namely, gap discontinuity, irregular gap discontinuity and orientation discontinuity. The measures for gap discontinuity, irregular gap discontinuity and orientation discontinuity of a patch P will be respectively denoted by $\mu_{GD}(P)$, $\mu_{IGD}(P)$ and $\mu_{OD}(P)$. The reader who is not interested in the technical details can jump directly to section 5.

In our analysis, we used chromatic range data of outdoor scenes[1] which were obtained from Riegl UK Ltd. (http://www.riegl.co.uk/). There were 20 scenes in total, 10 of which are shown in figure 3. The range of an object which does not reflect the laser beam back to the scanner or is out of the range of the scanner cannot be measured. These points are marked with blue in figure 3 and are not processed in our analysis. The resolution range of the data set is [512-2048]x[390-2290] with an average resolution of 1140x1001.

3D discontinuities are detected in studies which involve range data processing, using different methods and using different names like two-dimensional discontinuous edge, jump edge or depth discontinuity for gap discontinuity; and, two-dimensional corner edge, crease edge or surface discontinuity for orientation discontinuity [1, 8, 25].

[1]We would like to note that it is problematic to do range scanning in nature scenes that include trees or other kinds of vegetation because of the unintended motion due to wind. As the image of the scene is taken after the scanning phase, this delay may make the image data fail to correspond to the range data.

4.1. Measure for Gap Discontinuity: μ_{GD}

Gap discontinuities can be measured or detected in a similar way to edges in 2D images; edge detection processes RGB-coded 2D images while for a gap discontinuity, one needs to process XYZ-coded 2D images. In other words, gap discontinuities can be measured or detected by taking a second order derivative of XYZ values [25].

Measurement of a gap discontinuity is expected to operate on both the horizontal and vertical axes of the 2D image; that is, it should be a two dimensional function. The alternative is to discard the topology and do 'edge-detection' in sorted XYZ values, *i.e.*, to operate as a one-dimensional function. Although we are not aware of a systematic comparison of the alternatives, for our analysis and for our data, the topology-discarding gap discontinuity measurement produced better results. Therefore, we have adopted the topology-discarding gap discontinuity measurement in the rest of the paper.

For an image patch P of size $N \times N$, let,

$$\begin{aligned} \mathcal{X} &= \text{ascending\_sort}(\{X_i \mid i \in P\}), \\ \mathcal{Y} &= \text{ascending\_sort}(\{Y_i \mid i \in P\}), \\ \mathcal{Z} &= \text{ascending\_sort}(\{Z_i \mid i \in P\}), \end{aligned} \tag{1}$$

and also, for $i = 1, .., (N \times N - 2)$,

$$\begin{aligned} \mathcal{X}^{\Delta} &= \{ \mid (\mathcal{X}_{i+2} - \mathcal{X}_{i+1}) - (\mathcal{X}_{i+1} - \mathcal{X}_i) \mid \}, \\ \mathcal{Y}^{\Delta} &= \{ \mid (\mathcal{Y}_{i+2} - \mathcal{Y}_{i+1}) - (\mathcal{Y}_{i+1} - \mathcal{Y}_i) \mid \}, \\ \mathcal{Z}^{\Delta} &= \{ \mid (\mathcal{Z}_{i+2} - \mathcal{Z}_{i+1}) - (\mathcal{Z}_{i+1} - \mathcal{Z}_i) \mid \}, \end{aligned} \tag{2}$$

where $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{Z}_i$ represents 3D coordinates of pixel i.

The sets $\mathcal{X}^{\Delta}, \mathcal{Y}^{\Delta}$ and $\mathcal{Z}^{\Delta}$ are the measurements of the jumps (*i.e.*, second order differentials) in the sets $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$, respectively. A gap discontinuity can be defined simply as a measure of these jumps in these sets. In other words:

$$\mu_{GD}(P) = \frac{\phi(\mathcal{X}^{\Delta}) + \phi(\mathcal{Y}^{\Delta}) + \phi(\mathcal{Z}^{\Delta})}{3}, \tag{3}$$

where the function $\phi : \mathcal{S} \rightarrow [0, 1]$ over the set $\mathcal{S}$ measures the homogeneity of its argument set (in terms of its 'peakiness') and is defined as follows:

$$\phi(\mathcal{S}) = \frac{1}{\#(\mathcal{S})} \times \sum_{i \in \mathcal{S}} \frac{s_i}{\max(\mathcal{S})}, \tag{4}$$

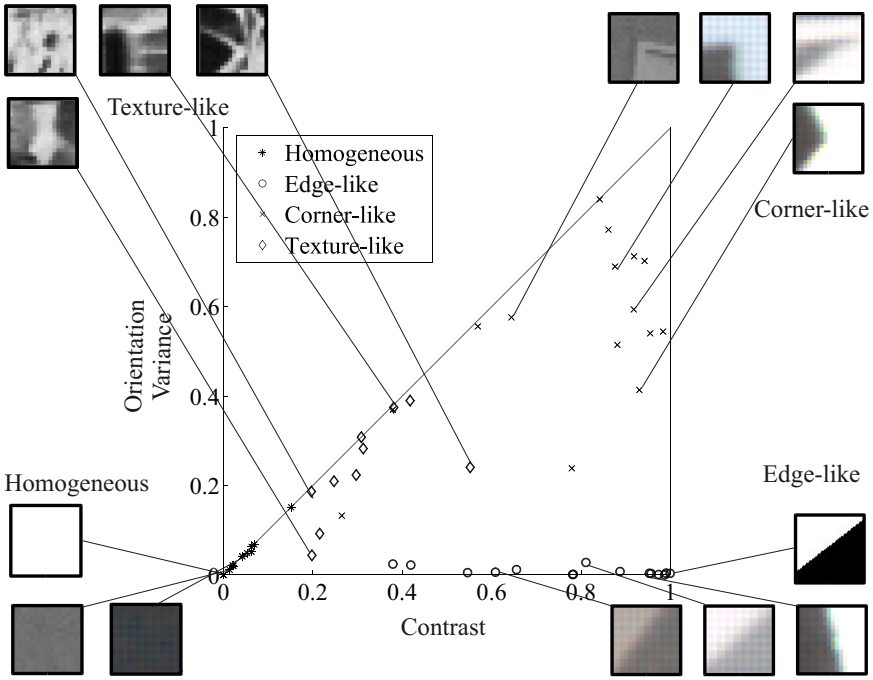

Figure 2. How a set of 54 patches map to the different areas of the intrinsic dimensionality triangle. Some examples from these patches are also shown. The horizontal and vertical axes of the triangle denote the contrast and the orientation variances of the image patches, respectively.

Figure 3. 10 of the 20 3D data sets used in the analysis. The points that don't have range data are marked in blue. The gray image shows the range data of the top-left scene. The resolution range is [512-2048]x[390-2290] with an average resolution of 1140x1001.

where $\#(\mathcal{S})$ is the number of the elements of $\mathcal{S}$, and s_i is the i^{th} element of the set $\mathcal{S}$. Note that as a homogeneous set (*i.e.*, a non-gap discontinuity) $\mathcal{S}$ produces a high $\phi(\mathcal{S})$ value, a gap discontinuity causes a low μ_{GD} value. Figure 5(c) shows the performance of μ_{GD} on one of our scenes shown in figure 3.

4.2. Measure for Orientation Discontinuity: μ_{OD}

The orientation discontinuity of a patch P can be detected or measured by taking the 3D orientation difference of the surfaces which meet at P. As the size of the patch P is small enough, the surfaces can be, in practice, approximated by 2-pixel wide unit planes. The histogram of the 3D orientation differences between every pair of unit planes forms one cluster for continuous surfaces and two clusters for orientation discontinuities.

For an image patch P of size $N \times N$ pixels, the orientation discontinuity measure is defined as:

$$\mu_{\mathbf{OD}}(P) = \psi(H^n(\{\alpha(i,j) \mid i,j \in planes(P), i \neq j\})), \tag{5}$$

where $H^n(S)$ is a function which computes the n-bin histogram of its argument set $\mathcal{S}$; $\psi(\mathcal{S})$ is a function which finds the number of clusters in $\mathcal{S}$; $planes(P)$ is a function which fits 2-pixel-wide unit planes to 1-pixel apart points in P using Singular Value Decomposition[2]; and, $\alpha(i,j)$ is the angle between planes i and j.

For a histogram H of size N_H, the number of clusters is:

$$\psi(S) = \frac{\sum_{i=1}^{N_H+1}(H_i > \frac{\max(H)}{10}) \neq (H_{i-1} > \frac{\max(H)}{10})}{2}, \tag{6}$$

[2] Singular Value Decomposition is a standard technique for fitting planes to a set of points. It finds the perfectly fitting plane if it exists; otherwise, it returns the least-square solution.

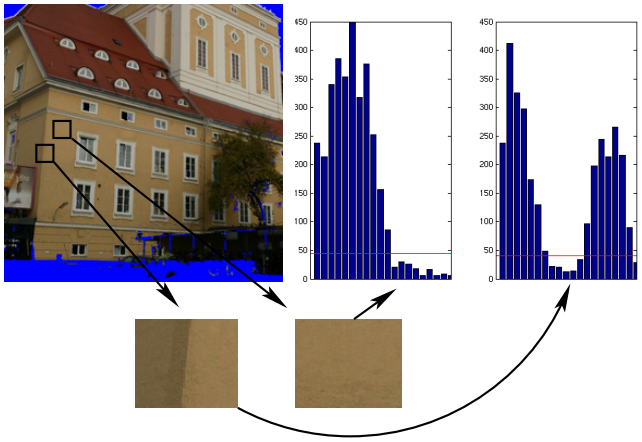

Figure 4. Example histograms and the number of clusters that the function $\psi(S)$ computes. $\psi(S)$ finds one cluster in the left histogram and two clusters in the right histogram. Red line marks the threshold value of the function. X axis denotes the values for 3D orientation differences.

where the operator $\neq$ returns 1 if its operands are not equal and returns 0, otherwise; H_i represents the i^{th} element of the histogram H; H_0 and H_{N_H+1} are defined as zero; and, $\max(H)/10$ is an empirical value which functions as the threshold value for finding the clusters. Figure 4 shows two example clusters for a continuous surface and an orientation discontinuity. Figure 5(d) shows the performance of μ_{OD} on one of our scenes shown in figure 3.

4.3. Measure for Irregular Gap Discontinuity: μ_{IGD}

Irregular gap discontinuity of a patch P can be measured by making use of the observation that an irregular-gap discontinuous patch from nature usually consists of small surface fragments with different 3D orientations. Therefore, the amount of variety in the 3D orientation histogram of a patch P can measure the irregular gap discontinuity of P.

Similar to the measure for orientation discontinuity defined in section 4.2, the histogram of the differences between the 3D orientations of the unit planes (which are of 2 pixels wide) is analyzed. For an image patch P of size $N \times N$ pixels, the irregular gap discontinuity measure is defined as:

$$\mu_{\mathbf{IGD}}(P) = \phi(H^n(\{\alpha(i,j) \mid i,j \in planes(P), i \neq j\})), \quad (7)$$

where $planes(P)$, $\alpha(i,j)$, $H^n(\mathcal{S})$ and $\phi(\mathcal{S})$ are as defined in section 4.2. Figure 5(e) shows the performance of μ_{IGD} on one of our scenes shown in figure 3.

The relation between the measurements and the types of the 3D discontinuities are outlined in table 1 which entails that an image patch P is:

- gap discontinuous if $\mu_{GD}(P) < T_g$ and $\mu_{IGD}(P) < T_{ig}$,
- irregular-gap discontinuous if $\mu_{GD}(P) < T_g$ and $\mu_{IGD}(P) > T_{ig}$,
- orientation discontinuous if $\mu_{GD}(P) \geq T_g$ and $\mu_{OD} > 1$,

Dis. Type	μ_{GD}	μ_{IGD}	μ_{OD}
Continuity	High value	Don't care	1
Gap Dis.	Low value	Low value	Don't care
Irregular Gap Dis.	Low value	High value	Don't care
Orientation Dis.	High value	Don't care	> 1

Table 1. The relation between the measurements and the types of the 3D discontinuities.

- continuous if $\mu_{GD}(P) \geq T_g$ and $\mu_{OD}(P) \leq 1$.

For our analysis, we have taken N and the threshold values T_g, T_{ig} empirically as 10, 0.4 and 0.6, respectively. The number of bins, n, in H^n is taken as 20.

Figure 5(a) shows the types of 3D discontinuities marked in four different colors for every pixel of the scenes shown in figure 3. We see that our measures can capture the 3D structure of the data sufficiently correct.

5. Results and Discussion

For each pixel of the scene (except for pixels where range data is not available), we computed the 3D discontinuity type and the intrinsic dimensionality. Figure 5(a) and (b) shows the images where the 3D discontinuity and the intrinsic dimensionality of each pixel are marked with different colors.

Having the 3D discontinuity type and the information about the local 2D structure of each point, it is straightforward to compute the probability $P(\text{3D Discontinuity} \mid \text{2D Structure})$, which is shown in figure 6. Note that the four triangles in figures 6(a), 6(b), 6(c) and 6(d) add up to one for all points of the triangle. We see that:

- Figure 6(a) shows that homogeneous image patches correspond to 3D continuities.

  Many surface reconstruction studies make use of a basic assumption that there is a smooth surface between any two points in the 3D world, if there is no contrast difference between these points in the image. This assumption has been first called as 'no news is good news' in [6]. With figure 6(a), we quantify 'no news is good news' and show for which structures and to what extent it holds. In addition to the fact that no news is in fact good news, the figure shows that news, especially texture-like structures and edge-like structures, can also be good news (see below).
- Edges are considered as important sources of information for object recognition and reliable correspondence finding. Approximately 10% of local image structures are of that type (see, *e.g.*, [12]). Figures 6(a), (b) and (d) show that most of the edges correspond to continuous surfaces or gap discontinuities. The edges that correspond to continuous surfaces are mostly low-contrast edges. Little percentage of the edges are formed by orientation discontinuities.

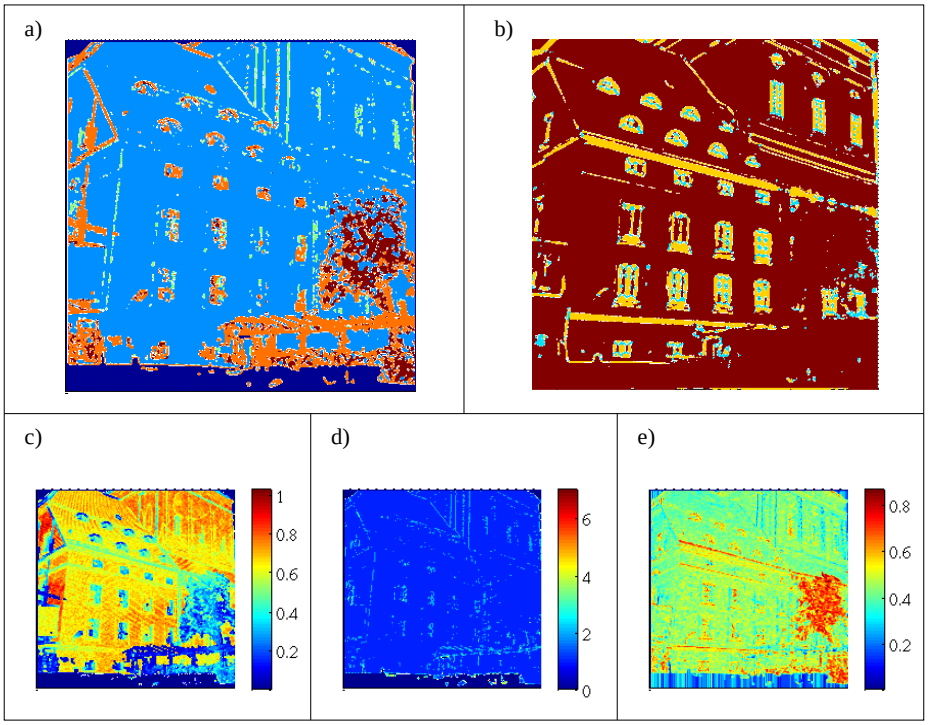

Figure 5. The 3D and 2D information for one of the scenes shown in figure 3. Dark blue marks the points without range data. (a) 3D discontinuity. Blue: continuous surfaces, light blue: orientation discontinuities, orange: gap discontinuities and brown: irregular gap discontinuities. (b) Intrinsic Dimensionality. Homogeneous patches, edge-like and corner-like structures are encoded in colors brown, yellow and light blue, respectively. (c) Gap discontinuity measure μ_{GD}. (d) Orientation discontinuity measure μ_{OD}. (e) Irregular gap discontinuity measure μ_{IGD}.

- Figure 6(b) shows that well-defined corner-like structures result from either gap discontinuities or continuities.
- Textures also map with high likelihood to surface continuities but also to irregular gap discontinuities.

  Finding correspondences becomes more difficult with the lack or repetitiveness of the local structure. The estimates of the correspondences at texture-like structures are naturally less reliable. In this sense, the likelihood that certain textures are caused by continuous surfaces (shown in figure 6(a)) can be used to model stereo matching functions that include interpolation as well as information about possible correspondences based on the local image information.

It is remarkable that local image structures mapping to different sub-regions in the triangle are caused by rather different 3D structures. This clearly indicates that these different image structures should be used in different ways for surface reconstruction.

6. Conclusion

In this paper, using 3D range data with real-world color information, we have analyzed the conditional probability of a 3D structure given the 2D structure. With this probability, we could investigate the relation between 2D structures and the underlying 3D structures as well as analyze the validity of a widely-used assumption/smoothing constraint, namely, 'no news is good news' [6].

Besides, we have presented a continuous classification scheme which can be used to distinguish between homogeneous, edge-like, corner-like and texture-like structures. By taking a higher-order representation than existing range-data analysis studies, we could point to the intrinsic properties of the 3D world and its relation to the image data. This analysis is important because (1) it may be that the human visual system is adapted to the statistics of the environment [2, 13, 15, 18, 21, 22], and (2) it may be used in several computer vision applications like depth estimation in a similar way as in [3, 4, 20, 29].

In our current work, the probability distributions will be used for estimating the 3D structure from 2D structure in a Bayesian framework for surface reconstruction/interpolation studies.

7. Acknowledgments

We would like to thank RIEGL UK Ltd. for providing us with 3D range data. This work is supported by the ECO-

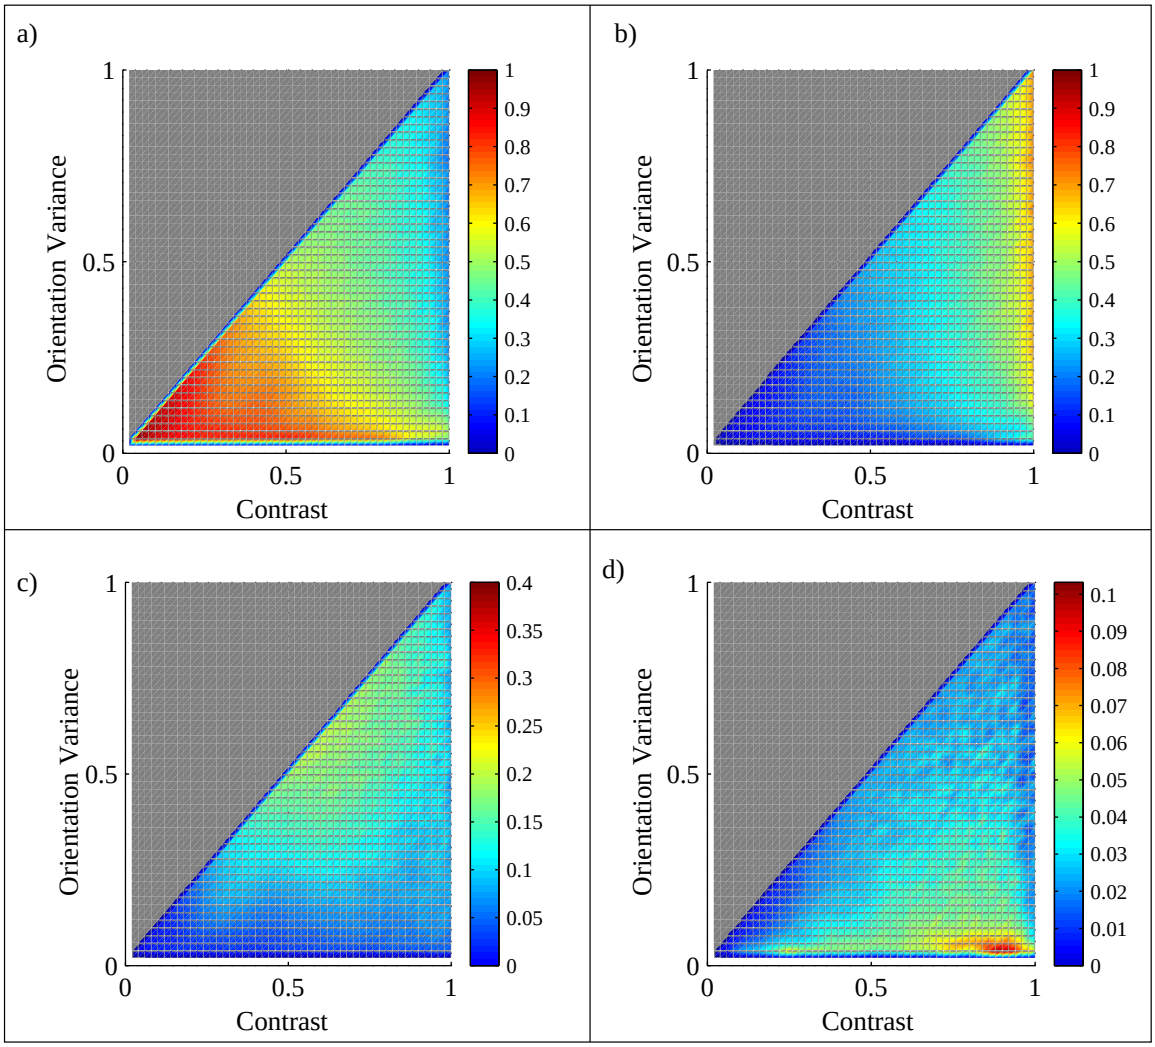

Figure 6. P(3D Discontinuity | 2D Structure): (a) P(Continuity | 2D Structure). (b) P(Gap Discontinuity | 2D Structure). (c) P(Irregular Gap Discontinuity | 2D Structure). (d) P(Orientation Discontinuity | 2D Structure).

VISION project.

References

[1] R. M. Bolle and B. C. Vemuri. On three-dimensional surface reconstruction methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(1):1–13, 1991. 3

[2] E. Brunswik and J. Kamiya. Ecological cue–validity of 'proximity' and of other Gestalt factors. *American Journal of Psychologie*, LXVI:20–32, 1953. 1, 6

[3] H. Elder and R. Goldberg. Ecological statistics of gestalt laws for the perceptual organization of contours. *Journal of Vision*, 2(4):324–353, 2002. 1, 6

[4] J. H. Elder, A. Krupnik, and L. A. Johnston. Contour grouping with prior models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(25):1–14, 2003. 1, 6

[5] M. Felsberg and N. Krüger. A probablistic definition of intrinsic dimensionality for images. *Pattern Recognition, 24th DAGM Symposium*, 2003. 2

[6] W. E. L. Grimson. Surface consistency constraints in vision. *Computer Vision, Graphics and Image Processing*, 24(1):28–51, Oct. 1983. 2, 5, 6

[7] A. Guzman. Decomposition of a visual scene into three-dimensional bodies. *AFIPS Fall Joint Conference Proceedings*, 33:291–304, 1968. 2

[8] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. B. Goldgof, K. Bowyer, D. W. Eggert, A. Fitzgibbon, and R. B. Fisher. An experimental comparison of range image segmentation algorithms.

IEEE Transactions on Pattern Analysis and Machine Intelligence, 18(7):673–689, 1996. 3

[9] C. Q. Howe and D. Purves. Range image statistics can explain the anomalous perception of length. *PNAS*, 99(20):13184–13188, 2002. 1

[10] C. Q. Howe and D. Purves. Size contrast and assimilation explained by the statistics of natural scene geometry. *Journal of Cognitive Neuroscience*, 16(1):90–102, 2004. 1

[11] J. Huang, A. B. Lee, and D. Mumford. Statistics of range images. *CVPR*, 1(1):1324, 2000. 1

[12] S. Kalkan, D. Calow, F. Wörgötter, M. Lappe, and N. Krüger. Local image structures and optic flow estimation. *Accepted for Network: Computation in Neural Systems*, 2005. 3, 5

[13] D. C. Knill and W. Richards, editors. *Perception as bayesian inference*. Cambridge: Cambridge University Press, 1996. 1, 6

[14] J. Koenderink and A. Dorn. The shape of smooth objects and the way contours end. *Perception*, 11:129—173, 1982. 2

[15] N. Krüger. Collinearity and parallelism are statistically significant second order relations of complex cell responses. *Neural Processing Letters*, 8(2):117–129, 1998. 1, 6

[16] N. Krüger and M. Felsberg. A continuous formulation of intrinsic dimension. *Proceedings of the British Machine Vision Conference*, 2003. 2

[17] D. Marr. *Vision: A computational investigation into the human representation and processing of visual information*. Feeman, 1977. 2

[18] B. Olshausen and D. Field. Natural image statistics and efficient coding. *Network*, 7:333–339, 1996. 1, 6

[19] B. Potetz and T. S. Lee. Statistical correlations between two-dimensional images and three-dimensional structures in natural scenes. *Journal of the Optical Society of America*, 20(7):1292–1303, 2003. 1

[20] N. Pugeault, N. Krüger, and F. Wörgötter. A non-local stereo similarity based on collinear groups. *Proceedings of the Fourth International ICSC Symposium on Engineering of Intelligent Systems*, 2004. 1, 6

[21] D. Purves and B. Lotto, editors. *Why we see what we do: an empirical theory of vision*. Sunderland, MA: Sinauer Associates, 2002. 1, 6

[22] R. P. N. Rao, B. A. Olshausen, and M. S. Lewicki, editors. *Probabilistic models of the brain*. MA: MIT Press, 2002. 1, 6

[23] N. Rubin. The role of junctions in surface completion and contour matching. *Perception*, 30:339–366, 2001. 2

[24] I. A. Shevelev, V. M. Kamenkovich, and G. A. Sharaev. The role of lines and corners of geometric figures in recognition performance. *Acta Neurobiol Exp*, 63(4):361–368, 2003. 2

[25] Y. Shirai. *Three-dimensional computer vision*. Springer-Verlag New York, Inc., 1987. 3

[26] M. Tuceryan and N. K. Jain. Texture analysis. *The Handbook of Pattern Recognition and Computer Vision (2nd Edition)*, pages 207–248, 1998. 2

[27] Z. Yang and D. Purves. Image/source statistics of surfaces in natural scenes. *Network: Computation in Neural Systems*, 14:371–390, 2003. 1

[28] C. Zetzsche and E. Barth. Fundamental limits of linear filters in the visual processing of two dimensional signals. *Vision Research*, 30, 1990. 2

[29] S. C. Zhu. Embedding gestalt laws in markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1170–1187, 1999. 1, 6