

Shape from texture: homogeneity revisited

A. Criminisi and A. Zisserman

Department of Engineering Science, University of Oxford

Parks Road, Oxford, OX1 3PJ, UK

[criminisi,az]@robots.ox.ac.uk

Abstract

The objective of this paper is to estimate the orientation of a scene plane from an uncalibrated perspective image under the assumption that the scene is coated with a homogeneous (but unknown) texture. We make the following novel contributions: first, we show that the problem is equivalent to estimating the vanishing line of the plane; second, we show that estimating the two degrees of freedom of this line can be decomposed into two searches each for one parameter; third, we give an algorithm for this estimation which is applicable to both regular and irregular textures. The algorithms do not require that texels are identified explicitly. But once the plane vanishing line has been obtained, then texels locations can be determined, and the geometry of the scene plane computed up to an affine transformation. We give examples of these computations on real images.

1 Introduction

In order to recover scene orientation from texture deformation in an image it is necessary to make assumptions, both on the camera (e.g. calibrated, orthographic) and on the nature of the texture (e.g. known, regular, homogeneous, isotropic). In this paper we analyze one of the most basic assumptions – *homogeneity*, for the case of a textured planar surface in a single, uncalibrated, perspective image. The objective is to remove the perspective deformation and estimate the orientation of the scene plane.

The homogeneity assumption allows only *projective*, not affine, properties of the imaged surface to be determined. In particular, for a planar surface only its *vanishing line* can be determined. This is because homogeneity is preserved under affine transformations (unlike *isotropy* or *compression* [1]). For example a regular square grid is transformed by an affine mapping to a grid of parallelograms all with the same shape and area, and constant density, i.e. to another homogeneous grid.

In terms of the imaging conditions perspective effects are necessary, otherwise the homogeneity assumption contributes no information on the vanishing line. If a plane is viewed by an affine camera (e.g. scaled orthography or weak perspective) then the imaged texture is homogeneous since there is an affine transformation between the world and image planes, and there is no relevant texture deformation that may be used to estimate the vanishing line.

Once the vanishing line of a plane has been computed the texture can be affine rectified. This affine rectification is very useful in many practical situations. For example: in single view metrology for measuring distances of points from planes [4]; in obtaining texture maps from images for graphical models [8]; in enabling a probabilistic analysis of slanted textures [11, 12]; and for texture synthesis algorithms which can not be applied to projective textures [5].

In contrast the texture assumption of isotropy (used by e.g. [2, 17]) allows the plane to be metric rectified such that a fronto-parallel view may be generated. However, this assumption is valid for fewer real world textures, and we do not investigate it further here.

In this paper the homogeneity assumption is “revisited” in the light of two developments since the work of [7, 11] a decade ago: (i) the projective geometry of uncalibrated images;

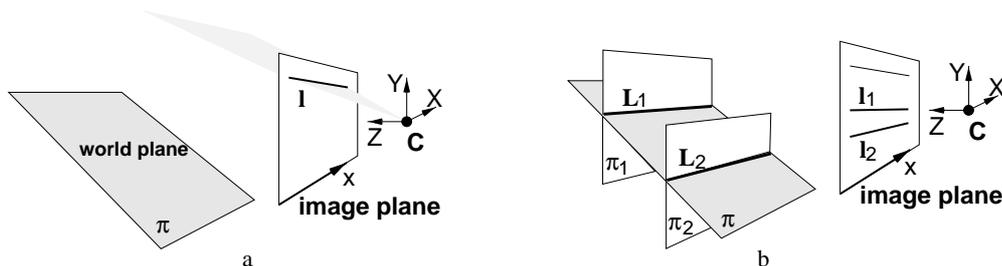


Figure 1: **Projective distortion:** (a) A camera (with centre C) images a world plane π parallel to the X -axis. The plane vanishing line l is parallel to the image x -axis. (b) The scene lines L_i are the intersection of π with planes parallel to the image plane, and are imaged at l_i . The sets of lines $\{L_i\}$, $\{l_i\}$, image x -axis, camera X -axis and vanishing line are parallel. The projective deformation along each image line l_i is constant (see text).

(ii) the use of the full texture deformation [10], not just the area deformation of [11]. The result is a new method of estimating plane orientation which does not require texels to be identified as a first step – a problem that has plagued many previous approaches.

2 Texture constraints on the vanishing line

In this section we describe the basic geometry of texture deformation and the induced within image mapping. The shape recovery algorithm, described in detail in section 3, is based on these results.

2.1 Constraints on the direction of the vanishing line

The key idea lies in the following observation:

Result 1 *In a perspective image of a plane, points on a line parallel to the vanishing line have the same perspective distortion.*

We first give a simple geometric proof. Suppose the scene plane π is parallel to the camera X -axis (fig. 1a) (this assumption is only to aid visualization). The vanishing line of the plane is then parallel to the image x -axis. Consider a set of fronto-parallel scene planes (i.e. planes parallel to the image plane of the camera, fig. 1b). Each of these planes intersects the plane π in a line parallel to the image plane. Consequently the perspective distortion along the image of this intersection line is the same (because points on the world line have the same depth). Each of these image lines is also parallel to the image x -axis and thus to the vanishing line.

From fig. 1 it is evident that on the retinal array the direction of the image gradient of perspective distortion is perpendicular to the vanishing line. However, a general calibration matrix (e.g. a non-unit aspect ratio) applies an affine transformation to the image. This transformation preserves parallelism, but not orthogonality. Thus in a general image the gradients of perspective distortion are parallel to each other, but their direction is not necessarily orthogonal to the vanishing line. Thus an important specialization of result 1 is

Result 2 *in the case of an image which has the correct aspect ratio (and has not been skewed) the gradient of perspective effects (corresponding to gradient in depth) is perpendicular to the vanishing line.*

The geometric results of this section capture Gibson's observation [6] that “texture gradient of the ground is orthogonal to the horizon on the retinal image” (see fig.2).

We now sketch a more formal proof for the case of imaged areas. In particular we show that the contour (lines) of imaged area deformation are parallel to the vanishing line.

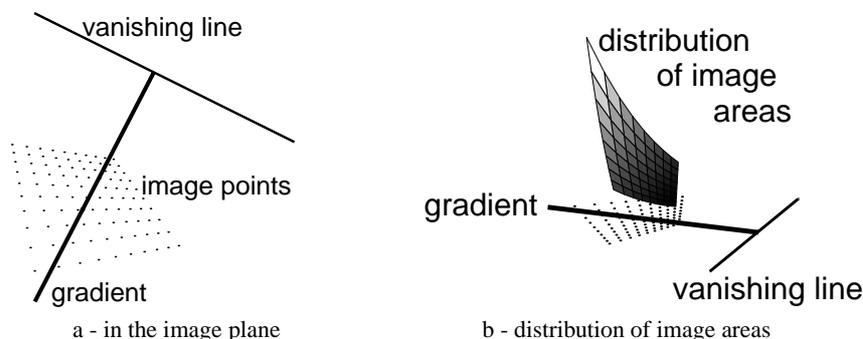


Figure 2: **The direction of the gradient of imaged area is orthogonal to the plane vanishing line:** (a) The (synthetic) perspective image of a regular square grid with vanishing line $\mathbf{l} = (0.005, 0.01, 1)^\top$. The gradient vector of the distribution of the areas of the imaged quadrilaterals is orthogonal to the vanishing line. (b) The 2D surface shows the area of the imaged quadrilaterals as a function of image position. The gradient of the surface is zero in the direction parallel to the vanishing line.

Proof. Consider a perspective image of a scene plane. Points on the scene and image plane are related by a homography matrix H , such that each world point \mathbf{X} maps to a point \mathbf{x} in the image as $\mathbf{x} = H\mathbf{X}$. Similarly lines are mapped as $\mathbf{l} = H^{-\top}\mathbf{L}$. Since we are only interested in projective properties of the image the matrix H may be chosen, without loss of generality, as

$$H = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ l_1 & l_2 & l_3 \end{bmatrix} \quad (1)$$

where $\mathbf{l} = (l_1, l_2, l_3)^\top$ is the vanishing line. This transformation maps the line at infinity of the scene plane $(0, 0, 1)^\top$ to \mathbf{l} .

If a world region \mathcal{R} undergoes a rigid translation on the scene plane, its image \mathcal{R}_i undergoes a more complicated transformation on the image plane – its shape changes according to image position. In particular the area of the imaged region $A_{\mathcal{R}_i}$ is a function of its image position \mathbf{x} and the world-to-image homography H : $A_{\mathcal{R}_i} = f(\mathbf{x}, H, A_{\mathcal{R}})$. At a point \mathbf{x} , $A_{\mathcal{R}_i} = A_{\mathcal{R}}|J|$, where $|J|$ is the Jacobian of $\mathbf{x} = H\mathbf{X}$ at \mathbf{x} . It can be shown that if the third coordinates of both \mathbf{x} and \mathbf{l} are 1 then $|J| = |H|(\mathbf{l} \cdot \mathbf{x})^{-3} = (\mathbf{l} \cdot \mathbf{x})^{-3}$.

Differentiating the Jacobian gives the area gradient $\left(\frac{\partial A}{\partial x}, \frac{\partial A}{\partial y}\right)^\top$ which is parallel to $(l_1, l_2)^\top$. Thus the level set consists of the lines $(l_1, l_2, d)^\top$ which are parallel to the vanishing line $(l_1, l_2, l_3)^\top$. \square

The length of the gradient vector, unlike its direction, is a function of the image point position. In particular the norm of the gradient is greater for points farther from the vanishing line.

2.2 Geometry given the vanishing line

A first result that we will use is that

Result 3 *if \mathbf{l} is identified in the image then the scene plane may be affine rectified using the transformation (1).*

The significance of this result is that this degree of rectification is all that is required in order to assess homogeneity on the (back-projected) scene plane.

To this point we have considered the homography between the scene and image plane. We now turn to the homography between the image and itself. If a pattern (e.g. a texel) repeats on a scene plane by a translation then the imaged pattern is related by a transformation which is conjugate to the world translation. This transformation is again a homography and is known as an *elation* [14, 15].

Elation. An elation is a plane-to-plane projective transformation which can be parametrized as $E = I + \mu \mathbf{v} \mathbf{a}^\top$, where \mathbf{v} is a fixed point corresponding to the translation direction; \mathbf{a} is a line of fixed points, and μ is a scalar related to the magnitude of the translation. The point \mathbf{v} is constrained to lie on the line \mathbf{a} ($\mathbf{v} \cdot \mathbf{a} = 0$). An elation has four dof: two for the line \mathbf{a} , one for the direction \mathbf{v} and one for the ratio μ . The full transformation can be computed from the correspondence $\mathbf{x}' = E\mathbf{x}$ of two pairs of points. In this case the line of fixed points is the vanishing line of the plane \mathbf{l} , so that $E = I + \mu \mathbf{v} \mathbf{l}^\top$.

Suppose for the moment that \mathbf{l} is known, then the two remaining degrees of freedom of the elation may be computed from a single point correspondence. In the case of a regular texture then, given \mathbf{l} , one correspondence between two texels determines the transformation between *all* texels related by repetition in the same direction.

3 Computing the vanishing line

This section describes our approach to estimating the vanishing line \mathbf{l} of the imaged textured plane.

The method is divided into two main steps: (i) estimating the direction of the vanishing line, and (ii) estimating its position. These steps stem from result 1. In the first step we search for the direction of a set of parallel lines in the image such that the perspective distortion is the same along each line (though differs between the lines). This determines one of the two degrees of freedom of the vanishing line. The second step then determines the remaining degree of freedom. Thus by employing result 1 we decompose the two-dimensional search for \mathbf{l} into two one-dimensional searches. We will consider the case of a texture which is *approximately* regularly repeating.

3.1 Step 1: estimating the direction of the vanishing line

We require a self-similarity *measure* on the image. Then we can determine the direction of minimum perspective distortion by finding the direction in which the self-similarity measure is a maximum.

The problem is thus reduced to one of measuring auto-correlation, and here we employ normalized auto-correlation for the similarity measure (NAC). Note, that NAC is invariant to an affine transformation of the intensities $I \rightarrow \alpha I + \beta$. This means that the measure is largely invariant to slowly varying illumination across the plane, where “slow” is in comparison to the region size used for the NAC.

Figure 3 shows an example of a perspectively distorted homogeneous texture and the NAC surface obtained by correlating a selected image patch with the entire image. This is a synthetic image generated by projectively distorting a fronto-parallel image of a real textured fabric. The region of the NAC surface surrounding the selected test patch is more corrugated – higher and sharper peaks, deeper valleys – than in other areas. The region orientation is horizontal in this case because the vanishing line is horizontal.

To quantitatively measure this effect we use the variance of the NAC values. Therefore, the direction of the vanishing line can simply be estimated by detecting the direction of the band for which maximum variance of the NAC values occurs. This is a one-dimensional search problem. The robustness of the algorithm is increased by repeating this search several

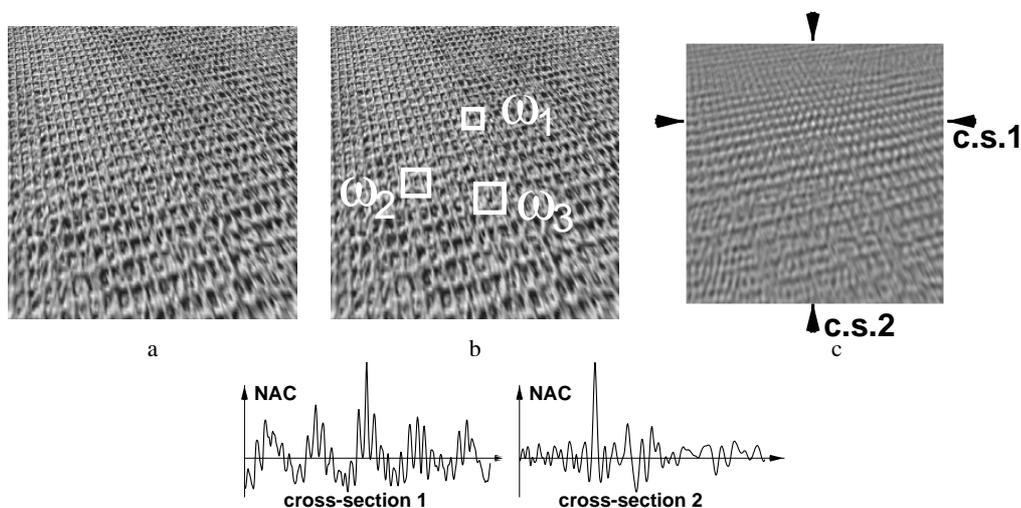


Figure 3: Peaks and valleys in the NAC surface: (a) A perspective image of plane fabric. The vanishing line is horizontal. The imaged texels have approximately equal area parallel to the vanishing line. Note – do not be distracted by the actual texture grid created by the fibres of the cloth but concentrate on the direction in which the texels maintain the same size. (b) Square patches ($\omega_1, \omega_2, \omega_3$, marked in white) of different sizes have been selected in the image. (c) The NAC surface related to one of the selected patches (ω_1). Notice the horizontal region where peaks and valleys of the NAC surface are sharper (sharper light and dark blobs). Two cross-sections of the NAC surface are also plotted: cross-section 1 is taken along a line parallel to the direction of minimum distortion (horizontal in this case) through the selected patch, ω_1 ; cross-section 2 is taken along a line parallel to the direction of maximum distortion (vertical in this case) through ω_1 . The curve corresponding to cross-section 1 is quite corrugated; it has a great variance of NAC values, along the whole line. Cross-section 2, instead, has “flatter” tails. The peak in the middle, corresponding to the strip centred on the selected patch, is ignored.

times for different patches positions and areas, and making a robust estimate from the set of computed directions (e.g. by using the median of directions, *cf.* fig. 4a). The algorithm is described in table 1.

The one-dimensional search in point (b) of the algorithm in table 1 may be implemented in different ways. Since we wish to find the global maximum of the variance of NAC values as a function of the orientation of the strip, attention must be paid to avoid local maxima (fig. 4a).

3.2 Step 2: estimating the position of the vanishing line

In this section we describe the second step of the algorithm to compute the complete vanishing line. It is, again, a 1D-search algorithm.

<ol style="list-style-type: none"> 1. Repeat <ol style="list-style-type: none"> (a) Select a random image patch (random centre \mathbf{x} and random area of the square correlation window ω) and compute the related NAC surface; (b) 1D search for the direction θ corresponding to the maximum value of variance of NAC values inside a strip of fixed width, centred on \mathbf{x} and with direction θ: <ul style="list-style-type: none"> • compute an initial value for θ (e.g. by maximizing NAC variance inside strip over a set of angles uniformly sampled in the range $[0^\circ, 180^\circ]$; <i>cf.</i> fig. 4a); • refine θ by non-linear maximization of the NAC variance (e.g. Levenberg-Marquardt). 2. The required direction corresponds to the median of the set of computed angles θ.

Table 1: Algorithm 1, estimating the direction of the vanishing line

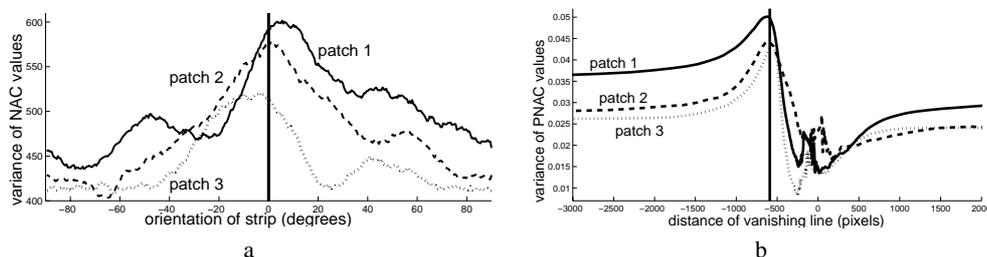


Figure 4: (a) **Variance of NAC values for rotating strips centred on different test points.** Each of the three curves shows the behaviour of the variance of the NAC values computed inside a strip of fixed size centred on a test point for directions varying from 0° to 180° . The three NAC surfaces have been computed for three different patches selected on the image in fig. 3a (horizontal vanishing line). The three surfaces correspond to the three patches marked in fig. 3b. The global maximum of each of the three curves is close to 0° (ground truth). (b) **Variance of PNAC values over entire image for translating vanishing line.** Each of the three curves shows the behaviour of the variance of the values of the whole PNAC surface for a vanishing line with fixed orientation and distance varying from $-\infty$ to $+\infty$. The three PNAC surfaces have been computed for three different patches. The vertical line in the diagram shows the distance corresponding to the ground truth vanishing line. This represents the situation of maximum homogeneity of the affine warped image. The global maximum of each of the three curves is close to the ground truth position. In both (a) and (b) the curves also show local maxima and minima.

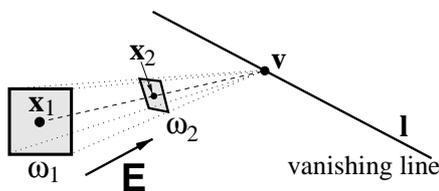


Figure 5: **Computing size and shape of correlation windows from elations:** Given an image point \mathbf{x}_1 , its correlation window ω_1 and the plane vanishing line l , the correlation window ω_2 for the point \mathbf{x}_2 is computed from the elation E of axis l and relating \mathbf{x}_1 with \mathbf{x}_2 . The vertex of the elation is the point $\mathbf{v} = (\mathbf{x}_1 \times \mathbf{x}_2) \times \mathbf{l}$ and $\mu = -\|\mathbf{x}_1 \times \mathbf{x}_2\| / \|(\mathbf{l} \cdot \mathbf{x}_1)(\mathbf{v} \times \mathbf{x}_2)\|$.

Having computed the direction of line l the only parameter that remains to be estimated is its position (say the distance of l from the centre of the image). We will build on result 3, i.e. that if the vanishing line is identified then the back-projected texture can be computed and its homogeneity assessed. The idea then is that if the plane is correctly rectified the texture is homogeneous in all directions, not just parallel to the vanishing line. Actually all computations are carried out in the image, to avoid sampling problems and having to tessellate to infinity on the scene plane.

We first describe a new, improved measure of similarity between any two image patches, which accounts for perspective distortion, and then show how to use it to compute the position of the vanishing line.

A new similarity function: projective correlation. Suppose we wish to measure the similarity between two points on the scene plane specified by image points \mathbf{x}_1 , \mathbf{x}_2 and an image correlation window ω_1 about \mathbf{x}_1 . If the vanishing line l is known an elation between the two points is completely constrained and therefore the size and shape of the corresponding correlation window ω_2 about \mathbf{x}_2 may be computed (fig. 5).

Once the window ω_2 is computed each point inside it can be warped by the inverse elation E^{-1} and the new image patch compared with the original (ω_1) by applying the usual normalized cross-correlation (this is possible because now the two patches have the same shape and size). This simple map provides the required invariance to perspective distortion. We have effectively defined a more powerful measure of similarity between two patches which is invariant to perspective distortion. The new similarity function is termed Projective

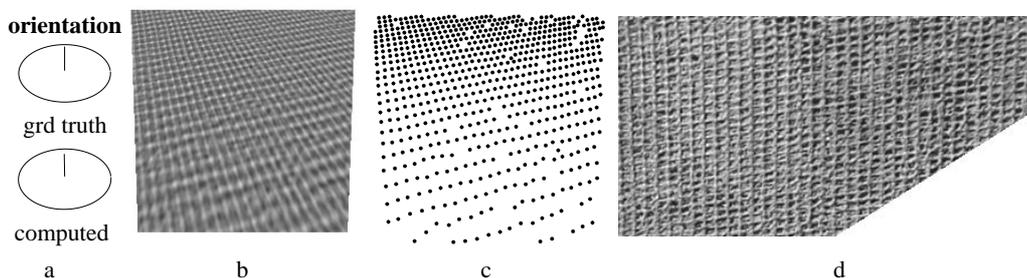


Figure 6: **Computing the vanishing line, the PNAC surface, rectifying the image and extracting texels, Ex. I:** (a) The vanishing line of the image of the plane fabric of fig. 3a has been computed by our algorithm (in sect. 3.1 and 3.2) to be: $\mathbf{l} = (5.04 \times 10^{-5}, 0.0027, 1)^T$. The ground truth is: $\mathbf{l} = (0, 0.003, 1)^T$. The two ellipses show the slight difference between computed and true orientations of the plane. (b) Once \mathbf{l} is known the PNAC image is computed for fig. 3a with respect to a selected test region. The peaks of the PNAC surface are now distributed, sharp and distinct over the entire image, no longer only inside a limited region (compare it with NAC image in fig. 3c). (c) Computed texel positions. Almost all the texels have been detected correctly in the image by thresholding the PNAC surface and grouping together the points obtained. (d) Affine rectification of the image in fig. 3a is obtained directly from the knowledge of \mathbf{l} (result 3). Notice that this is affine warping and therefore angles are not preserved but parallelism of lines is. The goodness of the estimate of the vanishing line can also be assessed by the degree of parallelism of the fibres in the rectified image.

NAC (PNAC).

Note: this map also provides a solution to the *scale selection* problem identified by [9, 13] (i.e. the computation of the correct size of neighbourhoods of image points such that the data inside the neighbourhood are representative of the texture characteristics), and the feature detection problem identified by [16] (i.e. in order to detect reliable image features it is necessary to compute a set of filters of different sizes for each image region). In the case of vanishing line at infinity then PNAC coincides with NAC.

Assessing a vanishing line. The scoring function is based on the deformation effect that the current vanishing line would have on the imaged texels. The score employed in this case is the variance of the values of the whole PNAC surface related to a selected test vanishing line (*cf.* fig. 4b). This has been found to be an effective method of measuring the homogeneity of the corrected image without explicitly correcting it.

Table 2 describes the algorithm for computing the remaining parameter of the vanishing line. The algorithm is analogous to that of table 1, except that the one-dimensional search is performed on a distance parameter rather than an angle.

3.3 Guided matching and texels extraction

Once an estimate of the vanishing line has been obtained a robust *guided* matching algorithm can be performed and the texels detected over the entire image.

Given a selected image region and the estimated vanishing line we can compute the *PNAC correlation surface* of the selected patch with every other point in the image (as earlier we computed the NAC surface, see fig. 6b). Since the similarity measure between image regions is no longer affected by perspective distortions but depends only on the world regions a simple thresholding and grouping algorithm on the PNAC surface is sufficient to detect the location of the image texels (fig. 6c).

MLE estimation of the vanishing line. Once the texel positions have been detected in the projective image, then, by applying the homogeneity assumption according to [11] (i.e.

<p>1. Repeat</p> <p>(a) Select a random patch (\mathbf{x}, ω) in the image.</p> <p>(b) 1D search of the second d.o.f. of the vanishing line (e.g. its signed distance d from the centre of the image; its orientation is fixed, computed as in table 1):</p> <ul style="list-style-type: none"> • compute an initial value for d (e.g. by maximizing variance of entire PNAC surface over a set of distances logarithmically sampled in the range $[-\infty, +\infty]$; cf. fig. 4b); • refine d by non-linear maximization of the variance of the whole PNAC surface (e.g. Levenberg-Marquardt). <p>2. The required vanishing line is the one corresponding to the median of the set of computed distances d.</p>

Table 2: Algorithm 2, estimating the second d.o.f. of the vanishing line

homogeneity = constant number of texels per unit area) it is possible to compute a ML Estimate of the plane vanishing line \mathbf{l} based on texels positions only.

The statistical theory developed in [11] is here reinterpreted in terms of projective geometry; rather than computing slant and tilt of a plane we compute its MLE vanishing line. It is straightforward to prove that in this projective framework the log-likelihood function of [11] becomes $\log L(\mathbf{l}) = -n \log(A_w) - 3 \sum_{i=1}^n \log |\mathbf{l} \cdot \mathbf{x}_i|$ where n is the number of texels in an image area A_i ; the corresponding world area is A_w and \mathbf{x}_i are the locations of each texel in the area A_i . Setting the derivatives of the log-likelihood function to zero yields (consistent with [11]) the result that the MLE of the plane vanishing line is obtained when the centroid of the back-projected area A_w coincides with the centre of mass of the back-projected points \mathbf{X}_i .

4 Examples

In this section we give results for: (i) the computed vanishing line \mathbf{l} ; (ii) the plane orientation; (iii) the affine rectified image; and (iv) the detected texels locations; on a set of example images. The plane orientation is computed as $\mathbf{n} = \mathbf{K}^T \mathbf{l}$ (where \mathbf{K} is the internal calibration matrix, and \mathbf{n} is the plane normal) and so requires an estimate of \mathbf{K} . However, the other three results do not require calibration.

The image in fig. 7a has been obtained by applying a known purely projective homographic warping to a fronto-parallel image from the set of textures in the VisTex database¹ (the procedure used for fig. 3a). This procedure provides a ground truth value for the vanishing line to which the computed vanishing line is compared. In figure 7b the orientation ellipses have been computed from the vanishing line by assuming a simple camera calibration; i.e. the principal point at the centre of the image and null skew. Figure 7d has been obtained by applying our rectification algorithm to fig. 7a. Notice in the rectified image the pattern elements have about the same area (ratios of areas are affine invariant). The accuracy of the algorithm is demonstrated by the closeness of the computed and ground truth orientation ellipses.

Figure 8 shows the results of the algorithm applied to a much less regular planar texture. The image in fig. 8a has been obtained by warping a fronto-parallel image (as in fig. 3a) (the original image is taken from a set of Brodatz [3] textures²). The algorithm accurately recovers the orientation of the plane and performs the rectification despite the weak regularities of the original pattern.

Exactly the same automatic algorithm works well also for very geometric texture patterns like the floor in fig. 9. In this image the “ground truth” vanishing line has been computed by

¹<http://www-white.media.mit.edu/vismod/imagery/VisionTexture/vistex.html>

²<http://www.ux.his.no/~tranden/brodatz.html>

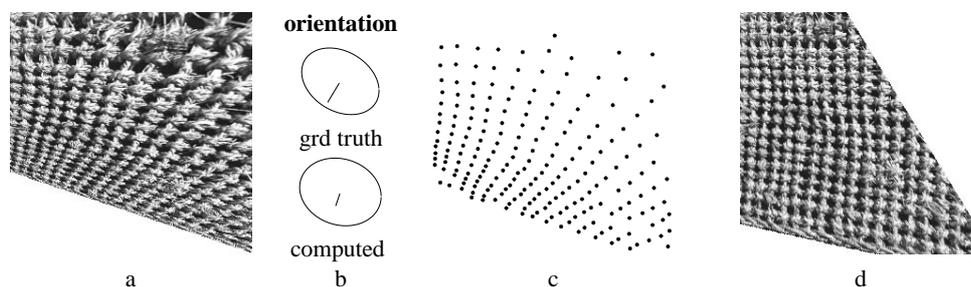


Figure 7: **Detecting texels and rectifying projective images, Ex. II:** (a) A perspective image of a textured planar surface. The true vanishing line is $\mathbf{l} = (0.001, -0.0016, 1)^T$. The computed vanishing line is $\mathbf{l} = (0.0005, -0.0013, 1)^T$. (b) Computed and ground truth orientations. (c) The detected texels positions (marked in black) using the computed vanishing line. (d) Image (a) affine rectified. In the rectified image angles are not preserved but areas are, in fact the basic elements of the cloth pattern all have about the same area.

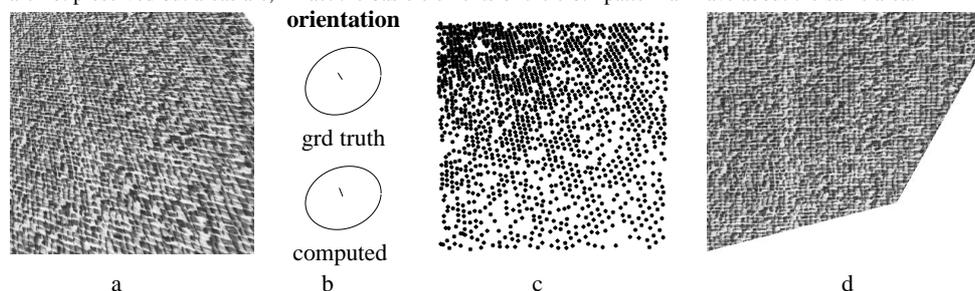


Figure 8: **Detecting texels and rectifying projective images, Ex. III:** Details as in fig. 7. The true vanishing line is $\mathbf{l} = (0.000769, 0.00123, 1)^T$ and the computed vanishing line is $\mathbf{l} = (0.000508, 0.00128249, 1)^T$.

manually defining and intersecting two sets of images of parallel world lines [8]. These have been defined by joining up sets of corresponding points in the tiles across the image.

Our algorithm is capable of automatically computing a very good estimate of the true vanishing line; this is demonstrated by comparison of the ground truth and computed orientation ellipses (fig. 9b) and the good parallelism of the tile pattern in the rectified image (fig. 9d).

If the assumption of homogeneity is not valid then the algorithm may fail already at its first step, the orientation of the vanishing line. A typical situation where the algorithm would show problems is in dealing with an image of a brick wall interrupted by windows and doors. The assumption of texture homogeneity is violated in the large by the interruptions in the brick pattern, and a robust analysis is required.

5 Conclusion and future work

A particular strength of the methods that we have developed is that, unlike previous algorithms [11, 12], it is not necessary to explicitly identify texels as the first step. The algorithm has been demonstrated for cases where texels are defined, and also for textures, such as figure 8, which are more stochastic. Even though the computation of the vanishing line does not need explicit detection of texels, once the vanishing line has been computed a robust guided-matching algorithm based on planar projective transformations may be performed to extract the texels locations.

We have set up a framework where different cost functions can be slotted in, and have investigated one such cost – the variance of the PNAC surface. It now remains to develop costs which are better tuned to stochastic textures. One such possibility is to sample and

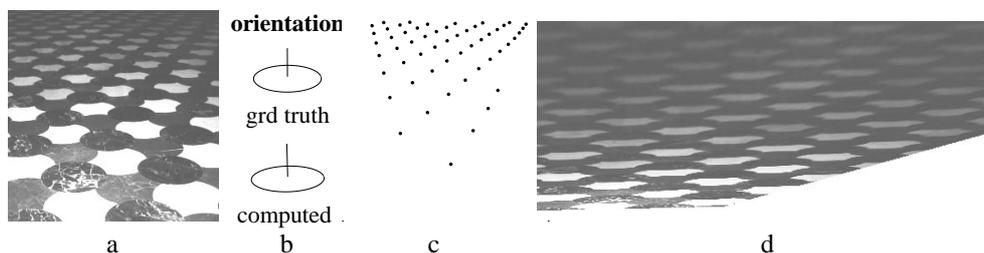


Figure 9: **Automatic rectification of slanted textured floors, Ex. IV:** (a) Original real image of a slanted, horizontal tiled floor. Its ground truth vanishing line is $\mathbf{l} = (0, 0.0077, 1)^T$. Notice the strong illumination gradient in this image. It is crucial that the similarity measures used in the algorithm are invariant to this gradient. (b) The computed vanishing line is $\mathbf{l} = (0.00024, 0.0093, 1)^T$. The computed orientation of the floor is compared to the ground truth. (c) The detected texels; the locations of the white tiles has been detected even in the dark region at the top. (d) Affine rectification of (a) by using the computed vanishing line. The parallelism of the tile pattern is recovered. Notice the strong blurring effect on the top tiles due to perspective reduction in imaged size.

synthesize in the manner of [5], which is a non-parametric method that does not require learning a texture pdf. A cost function may then be based on prediction error between the synthesized and measured texture.

Another extension is to use other attributes in the MLE. At present only the position of the texels is used, but with the PNAC cost function other attributes should be included, such as correlation strength.

Acknowledgements.

The authors would like to thank Prof. Andrew Blake, Frederik Schaffalitzky and Josephine Sullivan for discussions. This work was supported by the EU Esprit Project IMPROOFS.

References

- [1] A. Blake, H. H. Bülthoff, and D. Sheinberg. Shape from texture: Ideal observers and human psychophysics. *Vision Res.*, 33(12):1723–1737, 1993.
- [2] A. Blake and C. Marinos. Shape from texture: estimation, isotropy and moments. *Artificial Intelligence*, 1990.
- [3] P. Brodatz. *Textures: A Photographic Album for Artists & Designers*. Dover, New York, 1966.
- [4] A. Criminisi, I. Reid, and A. Zisserman. Single view metrology. In *Proc. ICCV*, pages 434–442, Sep 1999.
- [5] A. Efros and T. Leung. Texture synthesis by non-parametric sampling. In *Proc. ICCV*, pages 1039–1046, Sep 1999.
- [6] J. J. Gibson. *The Perception of the Visual World*. Houghton Mifflin, Boston, 1950.
- [7] K. Kanatani and T. Chou. Shape from texture: General principle. *AIJ*, 38:1–48, 1989.
- [8] D. Liebowitz, A. Criminisi, and A. Zisserman. Creating architectural models from images. In *Proc. Euro-Graphics*, volume 18, pages 39–50, Sep 1999.
- [9] T. Lindeberg and J. Gårding. Shape from texture from a multi-scale perspective. In *Proc. ICCV*, pages 683–691, May 1993.
- [10] J. Malik and R. Rosenholtz. Computing local surface orientation and shape from texture. *IJCV*, 2(23):143–168, 1997.
- [11] C. Marinos and A. Blake. Shape from texture: the homogeneity hypothesis. In *Proc. CVPR*, 1990.
- [12] Y. Ohta, K. Maenobu, and T. Sakai. Obtaining surface orientation from texels under perspective projection. In *Proc. 7th IJCAI*, pages 746–751, Vancouver, Canada, 1981.
- [13] E. Ribeiro and E.R. Hancock. Adapting spectral scale for shape from texture. In *Proc. ECCV*, LNCS 1842, pages 421–433. Springer-Verlag, 2000.
- [14] F. Schaffalitzky and A. Zisserman. Geometric grouping of repeated elements within images. In *Proc. BMVC*, pages 13–22, 1998.
- [15] J.G. Semple and G.T. Kneebone. *Algebraic Projective Geometry*. Oxford University Press, 1979.
- [16] J.V. Stone. Shape from texture: Textural invariance and the problem of scale in perspective images of textured surfaces. In *Proc. BMVC*. BMVA Press, 1990.
- [17] A. P. Witkin. Recovering surface shape and orientation from texture. *Artificial Intelligence*, 1981.