

Estimating Expansion Rates from Range Data Sequences

Hagen Spies
ICG-III: Phytosphere
Research Center Jülich
52425 Jülich, Germany
h.spies@fz-juelich.de

John L. Barron
Dept. of Computer Science
University of Western Ontario
London, Ontario, Canada, N6A 5B7
barron@csd.uwo.ca

Abstract

We present a method to compute surface expansion rates from sequences of range data. Towards this end the 3D velocity field (range flow) is extracted first and then used in a second step to estimate the local area expansion. A detailed performance analysis is presented and the method is applied on two real examples.

1 Introduction

We denote the instantaneous velocity field that describes the motion of a deformable surface as *range flow*. The term range flow is used as we derive this velocity field from sequences of range data sets. Together with the 3D structure the range flow field can be used to study the dynamic changes of such surfaces. One interesting question is whether the surface area changes during the motion. This can for example be used to study growth processes in biological systems such as leaves or skin.

The same displacement vector field has also been called scene flow when computed directly from stereo image sequences [14, 16, 3]. We present range flow estimation in a differential framework that is related to optical flow algorithms [8, 1, 9, 10, 11]. Other approaches that use deformable models have been reported before [12, 15].

Paper organisation: In Sect. 2 we review the concept of range flow estimation. Then we define a formula for local expansion rates in Sect. 3, an error analysis by means of error propagation is presented in Sect. 3.1. Experiments are reported in Sect. 4, where we first conduct a detailed performance analysis on synthetic test data (Sect. 4.1). Finally we apply our technique on two real examples in Sect. 4.2.

2 Range Flow

The concept of range flow estimation is briefly reviewed in the following. For a more detailed description we refer to [9, 10, 1].

2.1 Constraint Equations

Note that a range sensor produces one data set for each of X, Y and Z on its grid ($X = X(x, y, t)$ etc.). Here sensor coordinates are denoted by (x, y). The three components of the range flow field are the total derivatives of the world coordinates with respect to time ($U = \frac{dX}{dt}$ etc.). This can be expressed in the following equations:

$$U = \partial_x X \dot{x} + \partial_y X \dot{y} + \partial_t X, \quad (1)$$

$$V = \partial_x Y \dot{x} + \partial_y Y \dot{y} + \partial_t Y, \quad (2)$$

$$W = \partial_x Z \dot{x} + \partial_y Z \dot{y} + \partial_t Z. \quad (3)$$

The total time derivative is indicated by a dot. As we are not interested in the rates of change on the sensor coordinate frame we eliminate \dot{x} and \dot{y} to obtain the range flow motion constraint expressed in sensor coordinates:

$$\frac{\partial(Z, Y)}{\partial(x, y)} U + \frac{\partial(X, Z)}{\partial(x, y)} V + \frac{\partial(Y, X)}{\partial(x, y)} W + \frac{\partial(X, Y, Z)}{\partial(x, y, t)} = 0, \quad (4)$$

where $\frac{\partial(Z, Y)}{\partial(x, y)}$ is the Jacobian of Z, Y with respect to x, y. Notice that the Jacobians are readily computed from the derivatives of X, Y, Z in the sensor frame obtained by convolving the data sets with derivative kernels. In the following special derivative kernels optimised for directional invariance are used [7]:

$$\begin{aligned} \partial_x &= (0.084, 0.332, 0, -0.332, -0.084)_x \\ &* (0.023, 0.242, 0.470, 0.242, 0.023)_y \\ &* (0.023, 0.242, 0.470, 0.242, 0.023)_t. \end{aligned} \quad (5)$$

The other derivatives are computed in the same manner. In practice many sensors have aligned world and sensor coordinate systems which implies $\partial_y X = \partial_x Y = 0$. Yet Eq. (4) poses the general constraint independent of a particular sensor.

The usage of intensity data in addition to the range data can improve both the accuracy and density of the estimated range flow significantly [9]. We assume that the intensity

does not change with moderate depth changes. Thus, like for optical flow, we attribute all changes in intensity to motion. This yields another constraint equation:

$$0 = \partial_x I \dot{x} + \partial_y I \dot{y} + \partial_t I. \quad (6)$$

Combined with (1) and (2) we obtain an additional constraint on U and V :

$$\frac{\partial(I, Y)}{\partial(x, y)} U + \frac{\partial(X, I)}{\partial(x, y)} V + \frac{\partial(X, Y, I)}{\partial(x, y, t)} = 0. \quad (7)$$

2.2 TLS Solution

Equation (4) poses only one constraint in the three unknowns U, V, W . This manifestation of the aperture problem has been examined in more detail before [8]. In order to get an estimate we pool the constraints in a local neighbourhood and assume the flow to be constant within this area.

As all data terms in Eq. (4) are bound to contain errors it is reasonable to use total least squares estimation as opposed to standard least squares [13]. To do so we rewrite Eq. (4) as $\vec{d}^T \vec{p} = 0$ with a data vector \vec{d} given by the Jacobians. In order to avoid the trivial solution we require that $|\vec{p}| = 1$. It is straightforward to show that the solution is given by the eigenvector corresponding to the smallest eigenvalue of the so called structure tensor [8]:

$$J_{ij} = B * (d_i \cdot d_j), \quad i, j = 1 \dots 4. \quad (8)$$

The local integration is computed here via convolution with an averaging mask B , typically a Binomial. From the thus estimated parameter vector we can recover the range flow as (here \vec{p} is the eigenvector to the smallest eigenvalue):

$$\vec{f} = (U V W)^T = \frac{1}{p_4} (p_1 p_2 p_3)^T. \quad (9)$$

The smallest eigenvalue gives the residual of the attempted fit. This can be used to define a confidence measure based on a threshold τ [10]:

$$\omega = \begin{cases} 0 & \text{if } \lambda_4 > \tau \\ \left(\frac{\tau - \lambda_4}{\tau + \lambda_4} \right)^2 & \text{else} \end{cases}. \quad (10)$$

It is quite possible that the neighbourhood does not contain enough information to compute a full flow estimate. This can be somewhat amended by using the intensity data as well.

The intensity constraint (7) can also be written as $\vec{d}'^T \vec{p} = 0$, where we simply set $d'_3 = 0$. This results in another structure tensor J' constructed following Eq. (8). The sum of the two tensors yields a combined tensor from which the solution is then found by the analysis described above [9].

In order to ensure no a priori bias between intensity and depth we also require that the two data channels have been scaled such that their values are in the same order of magnitude. This can be done by subtracting the mean and adjusting the data to have the same variance. Additionally we can use a weight factor on the intensity tensor J' to account for different signal to noise ratios.

2.3 Regularisation

Even the use of the additional intensity data does not guarantee a full flow estimate everywhere. There are regions where the aperture problem prevails, sometimes the estimation fails and occasionally there is not enough variation in the data for a sensible flow calculation. We have previously presented a regularisation scheme that allows to close these remaining holes while taking the information from the TLS estimation into account [10].

We seek to estimate a dense and smooth flow field $\vec{v} = [U V W]^T$. To ensure smoothness we use a simply membrane model in this case, i.e. we require the first order changes of the flow field to be small. In places where flow estimations from the above TLS algorithm exist we denote them \vec{f} . The regularisation is then found as the solution of the following minimisation problem:

$$\int_A \left\{ \omega \left(P\vec{v} - \vec{f} \right)^2 + \alpha \sum_{i=1}^3 (\nabla v_i)^2 \right\} d\vec{r} \rightarrow \min. \quad (11)$$

Where ω , given by equation (10), captures the confidence of the TLS solution. At locations where no solution has been computed obviously no such data term exists. This is readily accounted for by setting the confidence values ω to zero. We use a projection matrix P which projects onto the subspace that was determined by our TLS algorithm. The overall smoothness can be regulated by the positive non-zero constant α . The minimum of (11) is achieved when the Euler-Lagrange equations are satisfied. The solution is readily obtained using an iterative update [10].

3 Surface Expansion

Let the considered surface \vec{s} be given by a set of range data parameterised in sensor coordinates x and y :

$$\vec{s} = \vec{s}(x, y) = [X(x, y) Y(x, y) Z(x, y)]^T. \quad (12)$$

We want to compute the local surface area and its change caused by a movement. The area near a surface point can be described by the area of the triangle created by moving a small distance h along the two parameter directions (Fig. 1):

$$a = \frac{1}{2} \left| [\vec{s}(x+h, y) - \vec{s}(x, y)] \times [\vec{s}(x, y+h) - \vec{s}(x, y)] \right|. \quad (13)$$




Figure 1: Computing the local surface area and its change caused by a non-uniform velocity.

Using the displacement vector field \vec{f} this area is transformed into (each point of the above triangle is moving):

$$a' = \frac{1}{2} \left| \left[\vec{s}(x+h,y) + \vec{f}(x+h,y) - (\vec{s}(x,y) + \vec{f}(x,y)) \right] \times \left[\vec{s}(x,y+h) + \vec{f}(x,y+h) - (\vec{s}(x,y) + \vec{f}(x,y)) \right] \right|. \quad (14)$$

The relative change in the local surface area is given by the quotient of a' and a . Taking the limit $h \rightarrow 0$ we obtain:

$$da = \lim_{h \rightarrow 0} \frac{\frac{1}{h} a'}{\frac{1}{h} a} = \frac{|\partial_x(\vec{s} + \vec{f}) \times \partial_y(\vec{s} + \vec{f})|}{|\partial_x \vec{s} \times \partial_y \vec{s}|}. \quad (15)$$

We can think of $a = |\partial_x \vec{s} \times \partial_y \vec{s}|$ as a local area element. The entire surface area of a regular surface can then be defined as the integral over these area elements [4]:

$$A = \int |\partial_x \vec{s}(x,y) \times \partial_y \vec{s}(x,y)| dx dy. \quad (16)$$

We are now ready to define the relative expansion rate as:

$$e = (da - 1) \cdot 100 \% = \frac{a' - a}{a} 100 \%. \quad (17)$$

This quantity can be computed from derivatives, obtained via convolution, of both the range data and range flow arrays. Both range flow and range data should be on the same scale, thus the same smoothing procedure is used on both data sets.

3.1 Error Analysis

The error in the expansion estimation caused by noise in the range data and range flow is estimated next. An experimental accuracy analysis using synthetic data sequences will be given in Sect. 4.1.

First recall that the error in the derivative is simply given by the error in the data multiplied by a factor γ depending on the used derivative kernel. For example for the derivatives in the depth map we obtain:

$$\sigma_{\partial_x Z} = \sigma_{\partial_y Z} = \gamma \sigma_Z \quad \text{with} \quad \gamma = \sqrt{\sum_{i=1}^n \alpha_i^2}, \quad (18)$$

where the α_i are the filter coefficients used. The 2D version of the optimised filters (5) result in a factor $\gamma = 0.28$. First the error in the local area element a is discussed.

In order to simplify the mathematics we make a few assumptions which are quite realistic for common range sensors. Because the lateral accuracy is typically about one order of magnitude higher than that in the depth only the noise in Z is considered [2]. We further assume no dependency of X on y and Y on x ($X_y = Y_x = 0$), this implies aligned sensor and world coordinate systems which is usually true to a good approximation. Let the coordinate system be oriented such that $Z_y = 0$ and assume the change in X and Y to be equal (square pixels): $X_x = Y_y = \Delta_{XY}$.

To clarify the influence of the various variables they are first treated separately. If the only change in the velocity field is in the X -component U in direction x ($U_x \neq 0$) then the error in the estimated expansion rate becomes:

$$\sigma_e = \frac{\gamma \sigma_Z}{\Delta_{XY}^2 (1 + k^2)^{\frac{3}{2}}} \times \frac{k U_x (U_x + 2\Delta_{XY})}{\sqrt{U_x^2 + 2U_x \Delta_{XY} + (1 + k^2) \Delta_{XY}^2}} 100 [\%]. \quad (19)$$

Here we used the slope $k = \frac{Z_x}{\Delta_{XY}}$ to express the error. Similar formulas are obtained for the other variables. Under the above assumptions a change of V in y does not result in an estimation error because there is no change in depth along the Y -direction. The error in the expansion rate as a function of the velocity gradient and the slope is shown in (Fig. 2), where $\Delta_{XY} = 0.185 \text{ mm/pixel}$ and $\sigma_Z = 0.1 \text{ mm}$ are used. The error depending on the horizontal velocity components increases with the surface slope and the velocity gradient (Fig. 2 a), the curves for U_y and V_x are very similar. For high slopes ($k > 0.5$) the standard deviation in the estimated expansion can be in the range of 5%. However, here and in the following one should bear in mind that the velocity gradients are seldom above 0.05. On a flat surface such a gradient corresponds to an expansion rate of 5%.




Figure 2: Error in the estimated area change in dependence of the velocity gradient and the depth slope.

A change in the vertical velocity component W yields an increasing error in the expansion rate even for a vanishing slope (Fig. 2 b). If we assume that there is no additional constant velocity a gradient in the vertical velocity corresponds to a rotation of the surface. Thus we conclude that a fast tilt of the observed surface can severely deteriorate the computed expansion rate. For an increasing slope in x the effect of the increasing vertical velocity is reduced (Fig. 2 b). As has to be expected the surface slope in x -direction has little influence on the expansion error caused by a change of W in y -direction.

The above considerations assumed that the error in the depth dominates the error in the expansion estimate. For high surface slopes this is true but for less inclined surfaces and small velocity gradients the error in the range flow field is an important factor. Assuming the same error σ_f in all three velocity components the error in the estimated expansion rate is directly proportional to the error in the range flow $\sigma_a = b \sigma_f$ with a factor b that depends on both the velocity gradient and the surface slope.

Varying only one velocity gradient component and the surface slope at a time results in relatively flat curves for this factor (Fig. 3). There is the general trend that higher velocity gradients result in higher expansion errors and that an increasing slope reduces the error. Only for the change of the vertical velocity component in direction of the surface orientation W_x (Fig. 3 c) we find a slight increase with the slope for higher velocity gradients. Overall the deviation in the estimated expansion, caused by the error in the velocity alone, is given by about 2 times the deviation in the velocity itself. An error of $\sigma_f = 0.01 \text{ mm/h}$ in the velocity leads to a standard deviation in the expansion of $2\%/h$.

Such a high variation is caused by the fact that expansion is in fact a second order process. The absolute error in a derivative is about one third that in the data. But as the derivative values are at least a factor 10 smaller than the data values this increases the relative error by a factor 3.

Finally the combined error caused from both uncertain depth and uncertain velocities is examined qualitatively. As the error propagation formulas become very lengthy they are




Figure 3: Factor relating the error in the range flow estimation to the error in the expansion rate, depending on a velocity gradient and the surface slope ($V_x = V_y = 0$). **a** U_x , **b** U_y , **c** W_x and **d** W_y .

omitted here. The expansion error for two example velocity errors $\sigma_f = 0.01$ and $\sigma_f = 0.02$ (Fig. 4) show that for higher velocity noise levels the error is dominated by the velocity uncertainty. In this regime the expansion estimation is not very accurate and only the integration over a large area can reduce the error. For lower noise in the velocity the error is dominated by the velocity for small slopes and small velocity gradients. But for high velocity and depth gradients the error in the depth dominates the expansion estimation. Thus it becomes unreliable for higher slopes.

3.2 Reducing the Error by Averaging

For real range data the relative error remains below 1% only as long as the movement does not exceed a speed of 0.5 mm/frame on structured light data. At a sampling rate of 3 minutes this implies an absolute error in the velocity less than $\sigma_f = 0.01 \text{ mm/h}$. The resulting maximal uncertainty in the expansion estimate thus becomes $\sigma_e = 2\%/h$. If we want to reduce this error we need to compute the expansion on a coarser spatial scale than the velocity itself.

Under the assumption that the expansion rate does not change rapidly averaging will reduce the statistical error. Such an averaging can be done very efficiently using a Gaussian pyramid [6]. Here the Pyramid is realised by applying a 5-tap Binomial pre-filter then sub-sampling and finally applying a 3-tap Binomial post-filter [5]. The resulting effective 1D mask is given by: $\frac{1}{64}(1, 4, 8, 12, 14, 12, 8, 4, 1)$. This filter is applied first in x and then in y -direction yielding an overall noise suppression factor $\gamma = 0.158$. Each




Figure 4: Relative error in the estimated surface expansion in dependence of the depth slope k and the velocity gradient (U_x , W_x and W_y) for two velocity noise levels: $\sigma_f = 0.01$ (**a**, **b**, **c**) and $\sigma_f = 0.1$ (**d**, **e**, **f**). The error in the depth data is assumed to be $\sigma_z = 0.1$.

level of the pyramid reduces the noise by this factor. On the second level the uncertainty in the expansion estimate thus reduces to $\sigma_e = 0.025 \cdot 2\%/h = 0.05\%/h$ which is quite acceptable.

4 Experiments

This section presents results of the proposed expansion estimation framework. First a detailed performance analysis is done using synthetic data, then we present some results on real data in Sect. 4.2.

4.1 Synthetic Data

To investigate the performance on realistic test data we model a sensor using perspective projection, which corresponds to all sensors that employ a standard camera. Here we use a focal length of $f = 12\text{ mm}$, a pixel size of $7.4 \times 7.4 \mu\text{m}^2$ and 256×256 sensor elements. For testing a planar target as well as a sphere are used.

For the synthetic plane we choose a viewing distance of




Figure 5: Synthetic sphere: **a** texture mapped range data, **b** intensity data and **c** depth (Z).

300 mm and typically an angle between the surface normal and the Z -axis of $\vartheta = 5^\circ$. In this case a sinusoidal plaid pattern is used as intensity. The second type of synthetic data used consists of a sphere with a radius of 300 mm with its centre initially placed 700 mm away from the camera. Here we apply an intensity texture to each surface element based on the spherical angles $I = I(\theta, \phi)$:

$$I = \begin{cases} 100 & \text{if } \theta < 0.5^\circ \\ 100 + 50 \sin\left(\frac{2\pi\theta}{1^\circ}\right) + 50 \sin\left(\frac{2\pi\phi}{30^\circ}\right) & \text{else} \end{cases} \quad (20)$$

An example of the thus created data is given in Fig. 5. Note that in order to obtain the correct range flow at each sensor location the velocity has to be projected into the sensor frame if a non-rigid motion is used.

For the expansion estimation both range flow and range data are required to be available on the same scale. Thus we also smooth the range data using the membrane model ($\alpha = 10$) described in Sect. 2.3. Note that the smoothing is only done for the expansion rate estimation and not for the range flow computation. This is because on real data such a regularisation step would reduce much of the needed variation in the data.

In the following we will examine the performance for three typical noise situations:

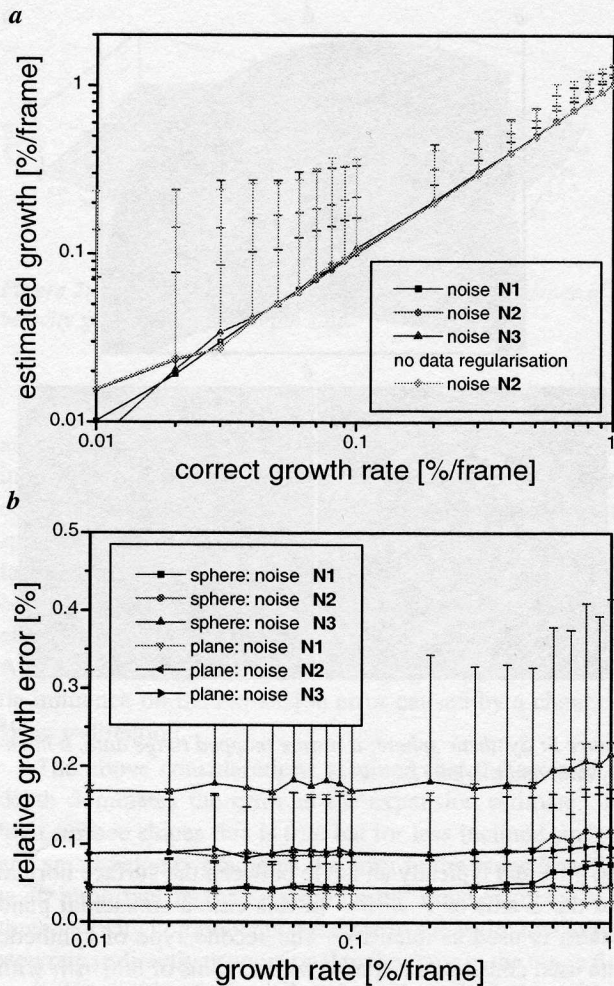


Figure 6: Expansion rate estimation: *a* estimated versus correct expansion rate on the expanding sphere and *b* relative error in the expansion rate for increasing correct expansion on the sphere and plane data.

| name | $\sigma_X = \sigma_Y$ | σ_Z | σ_I |
|------|-----------------------|------------|------------|
| N1 | 0.005 | 0.05 | 0.5 |
| N2 | 0.01 | 0.1 | 1.0 |
| N3 | 0.02 | 0.2 | 2.0 |

Independent normally distributed noise with these standard deviations is added to the X, Y, Z, I data arrays.

The mean computed expansion rate and its standard deviation as a function of an increasing correct expansion rate on the sphere data are shown in Fig. 6 a. On average the expansion rate is computed accurately from 0.01 %/frame up to 1 %/frame, for even higher expansion rates the estimation quickly fails completely. However, there is a substantial spread in the result. This implies that a lot of averaging is needed in this case to achieve reliable results. The variance in the computed expansion increases with the noise level and

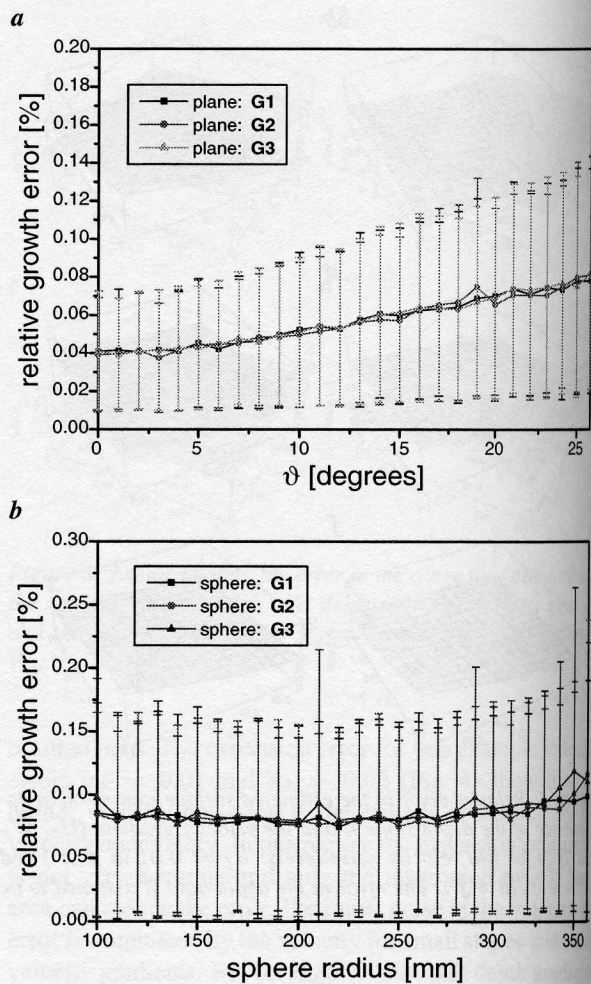


Figure 7: Error in the estimated expansion rate: *a* depending on the orientation of the plane ϑ and *b* on the initial radius of the expanding sphere.

we find a rather big difference between N1 and N3. Note that the effect of the noise is already reduced by the regularisation of the data, without it we find an even higher variance in the computed expansion rate (Fig. 6 a).

To quantify the accuracy of the calculated expansion rate we report the mean relative error. Let the correct expansion be given by e_c and the estimated rate by e_e , then the *relative expansion error* becomes:

$$E_r = \frac{|(|e_c| - |e_e|)|}{|e_c|} \cdot 100 [\%]. \quad (21)$$

The resulting mean relative error in dependence of the correct expansion rate is shown in Fig. 6 b. First we note that the error increases with the noise level as has to be expected. There is only a minor increase of the mean error for higher expansion rates. The error is somewhat smaller on the planar data than it is on the sphere. Yet in both cases the estimated

expansion rate agrees very well with the correct rate.

In the following we examine the influence of additional movements on the estimated expansion rate. This will be done on the realistic noise regime **N2**. Furthermore we only investigate a few exemplary expansion (growth) rates:

| name | expansion rate [%/frame] |
|-----------|--------------------------|
| G1 | 0.025 |
| G2 | 0.1 |
| G3 | 0.25 |

The error for these expansion rates on the plane in dependence of the vertical angle ϑ is given in Fig. 7 a. There is a small increase in the error for larger angles but the error is very small even then. On the sphere we explore whether the radius and thus the curvature has an effect on the expansion estimation. The relative error as a function of the radius is given in Fig. 7 b. No dependency on the radius can be observed. Also note that in both cases the value of the correct expansion rate has no effect on the achieved accuracy, in the considered cases. Thus we will only look at a rate of 0.1 %/frame (**G2**) in the following.

Next we analyse the influence of an additional movement on the accuracy of the estimated expansion rate. The error in the expansion depending on an extra translation in the X- or Z-direction is shown in Fig. 8 a for the plane at $\vartheta = 5^\circ$ and the sphere with radius $r = 300\text{ mm}$. For both data sets we find better results for a motion along the Z-axis (W). This is obviously caused by the more accurate velocity estimation for such movements. In the same way as the velocity estimation fails for movements above 1 mm/frame in X-direction does the expansion estimation become unreliable. This failure is caused by aliasing and of course dependent on the sensor and scene geometry.

On the synthetic plane data we can investigate whether the determined expansion rate is influenced by a rotational movement. Figure 8 b gives the error in the computed expansion rate when there is an additional change in either ϑ or ϕ . A rotation around the Z-axis, i.e. a change in ϕ , has no effect for rotations of up to $10^\circ/\text{frame}$. This is not the case for a rotation increasing the angle ϑ between Z-axis and surface normal. Here we find that a change by more than $1^\circ/\text{frame}$ leads to a large error in the expansion rate.

4.2 Real Data

Figure 9 shows two real examples and the thus computed expansion maps. The first example consists of a sequence of a yeast dough taken with a Biris range sensor at a sampling rate of one minute. One example depth map is given in Fig. 9 a. The computed range flow field (Fig. 9 b and c) shows a diverging X, Y velocity and a movement in Z that is higher on top of the object. This range flow is part of the

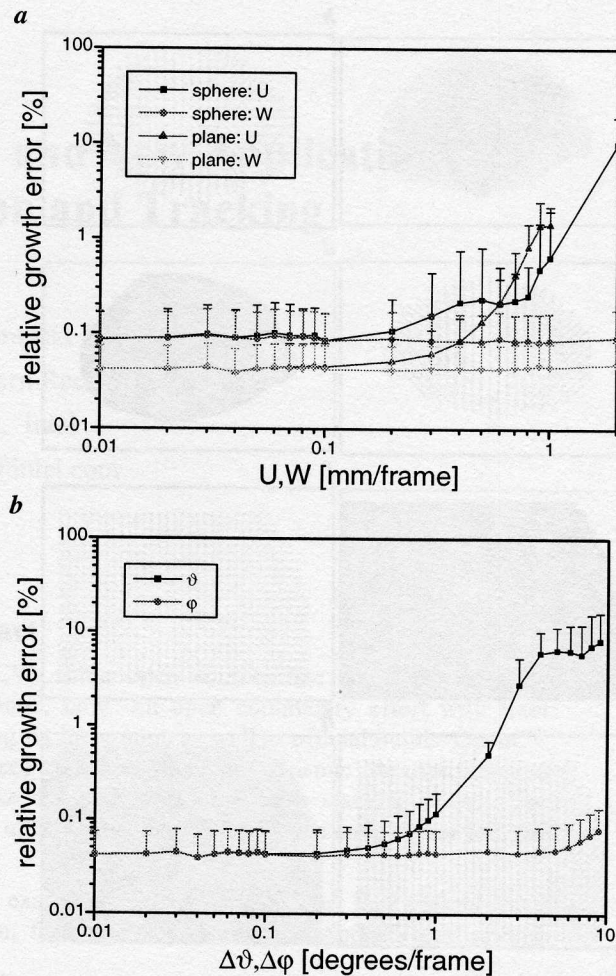


Figure 8: Error in the estimated expansion rate of 0.1 %/frame depending on additional movement: **a** translation in X- or Z-direction and **b** rotation (ϕ) and tilt (ϑ) of the plane.

3D diverging velocity caused by an increase in volume. On most parts of the dough we find an areal expansion around $20\%/h$ (Fig. 9 d). There are some places showing an even higher expansion mainly at locations with a steep depth gradient. Even though the expansion might vary in the observed manner the extracted expansion rate is probably not as reliable at these locations, see Sect. 3.1.

For the second example an air balloon is measured using a structured light system, see Fig. 9 e for an example depth map. During the observation time the air slowly vanishes and the balloon shrinks, scans are taken at one minute sampling intervals. This results in a velocity field with negative divergence in X, Y and an overall downward movement (Fig. 9 f and g). The expansion map (Fig. 9 h) shows a shrinkage of about $-3\%/h$. Again the expansion map is more noisy in regions with a steeper depth, as noted above.

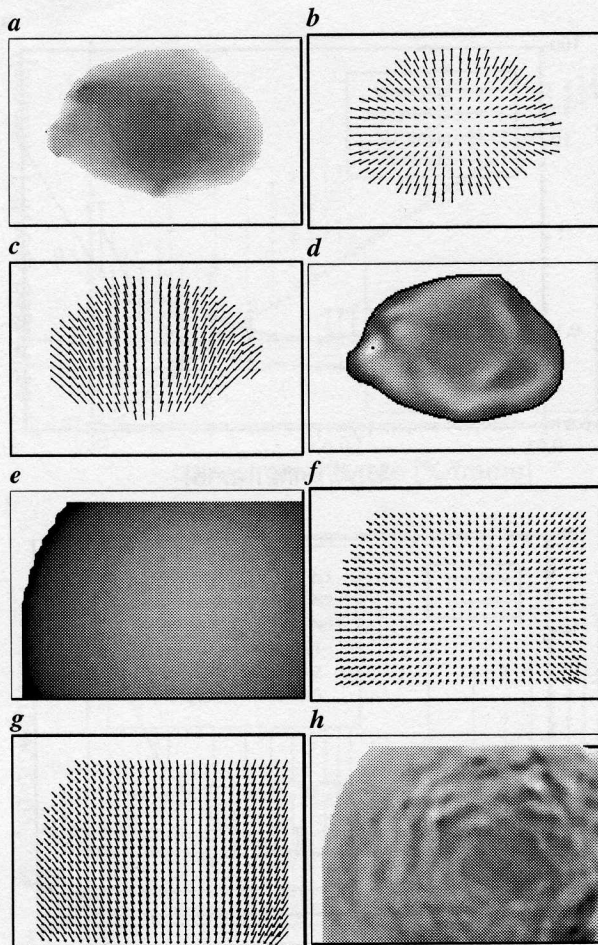


Figure 9: Example expansion map on real sequences. An expanding yeast dough: **a** depth data, **b** XY-components of the range flow, **c** XZ-components of the range flow field and **d** expansion map in the range of $[0, 40]\%/h$. A shrinking air balloon: **e** depth data, **f** XY-components of the range flow, **g** XZ-components of the range flow field and **h** expansion map in the range of $[-6, 0]\%/h$.

5 Conclusion

A method to compute local surface expansion rates on deformable surfaces has been presented. Towards this end we reviewed range flow estimation and presented an algorithm to compute the expansion rate from the thus available dense velocity field. A detailed error analysis by means of both error propagation and tests on synthetic data has been presented. We also should that the method has been successfully applied on real data.

Acknowledgments

Part of this work has been funded under the DFG research unit "Image Sequence Analysis to Investigate Dynamic Processes"

(FOR240) and an NSERC operating grant.

References

- [1] J. L. Barron and H. Spies. Quantitative regularized range flow. In *Vision Interface*, pages 203–210, Montreal, Canada, May 2000.
- [2] J. Beraldin, S. F. El-Hakim, and F. Blais. Performance evaluation of three active vision systems built at the national research council of Canada. In *Conf. on Optical 3D Measurement Techniques III*, pages 352–361, Vienna, 1995.
- [3] R. Carceroni and K. Kutulakos. Multi-view scene capture by surfel sampling: From video streams to non-rigid 3d motion, shape & reflectance. In *ICCV*, pages 60–67, Vancouver, Canada, July 2001.
- [4] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976.
- [5] H. Haußecker and B. Jähne. Ein Mehrgitterverfahren zur Bewegungssegmentierung in Bildfolgen. In *DAGM*, pages 24–31, 1993.
- [6] B. Jähne. *Digital Image Processing*. Springer, Berlin, Germany, 3 edition, 1995.
- [7] B. Jähne, H. Scharr, and S. Körkel. Principles of filter design. In B. Jähne, H. Haußecker, and P. Geißler, editors, *Handbook of Computer Vision and Applications*, volume 2, pages 125–151. Academic Press, 1999.
- [8] H. Spies, H. Haußecker, B. Jähne, and J. L. Barron. Differential range flow estimation. In *DAGM*, pages 309–316, Bonn, Germany, September 1999.
- [9] H. Spies, B. Jähne, and J. L. Barron. Dense range flow from depth and intensity data. In *ICPR*, pages 131–134, Barcelona, Spain, September 2000.
- [10] H. Spies, B. Jähne, and J. L. Barron. Regularised range flow. In D. Vernon, editor, *ECCV*, volume 2 of *Lecture Notes in Computer Science 1843*, pages 785–799, Dublin, Ireland, June/July 2000. Springer.
- [11] H. Spies, B. Jähne, and J. L. Barron. Surface expansion from range data sequences. In *DAGM*, Lecture Notes in Computer Science 2191, pages 163–169, Munich, Germany, September 2001. Springer.
- [12] L. Tsap, D. Goldgof, and S. Sarkar. Multiscale combination of physically-based registration and deformation modeling. In *CVPR*, volume 2, pages 422–429, June 2000.
- [13] S. Van Huffel and J. Vandewalle. *The Total Least Squares Problem: Computational Aspects and Analysis*. Society for Industrial and Applied Mathematics, Philadelphia, 1991.
- [14] S. Vedula, S. Baker, P. Rander, R. Collins, and T. Kanade. Three-dimensional scene flow. In *ICCV*, volume 2, pages 722–729, Corfu, Greece, September 1999.
- [15] M. Yamamoto, P. Boulanger, J. Beraldin, and M. Rioux. Direct estimation of range flow on deformable shape from a video rate range camera. *PAMI*, 15(1):82–89, January 1993.
- [16] Y. Zhang and C. Kambhamettu. Integrated 3d scene flow and structure recovery from multiview image sequences. In *CVPR*, volume 2, 2000.