Surface Expansion from Range Data Sequences

Hagen Spies[1], Bernd Jähne[1], and John L. Barron[2]

[1] Interdisciplinary Center for Scientific Computing,
University of Heidelberg, INF 368, 69120 Heidelberg, Germany,
{Hagen.Spies,Bernd.Jaehne}@iwr.uni-heidelberg.de
[2] Dept. of Comp. Science, University of Western Ontario,
London, Ontario, N6A 5B7 Canada,
barron@csd.uwo.ca

Abstract. We compute the range flow field, i.e. the 3D velocity field, of a moving deformable surface from a sequence of range data. This is done in a differential framework for which we derive a new constraint equation that can be evaluated directly on the sensor data grid. It is shown how 3D structure and intensity information can be used together in the estimation process. We then introduce a method to compute surface expansion rates from the now available velocity field. The accuracy of the proposed scheme is assessed on a synthetic data set. Finally we apply the algorithm to study 3D leaf motion and growth on a real range sequence.
Keywords. *range flow, expansion rates, range data sequences.*

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 [11]. We present range flow estimation in a differential framework that is related to optical flow algorithms. Other approaches that use deformable models have been reported before [10, 12].

The contribution of this paper is twofold. First we introduce a new version of the constraint equation for range flow estimation that can be evaluated directly on the sensor grid. Second we show how a change in the surface area can be determined locally from range flow fields.

2 Range Flow

We now restate the concept of range flow estimation and introduce a new formulation of the range flow motion constraint equation.

2.1 Motion Constraint

The observed surface can be described by its depth as a function of space and time $Z = Z(X, Y, t)$. The total derivative with respect to time then directly yields the range flow motion constraint equation [4, 12, 9]:

$$Z_X U + Z_Y V - W + Z_t = 0 \ . \tag{1}$$

Here $(U, V, W)^T$ is the range flow and indices denote partial derivatives. In order to evaluate this equation we need to compute these partial derivatives of the depth function with respect to world coordinates. However typical range sensors sample the data on a sensor grid, e.g. that of a CCD camera.

While it is possible to compute derivatives from a surface fit, we rather use convolutions. This allows to draw on the well established linear filter theory and can be implemented very efficiently. Towards this end we notice 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 ($U = \frac{\mathrm{d}X}{\mathrm{d}t}$ etc.). This can be expressed in the following equations:

$$U = \partial_x X \, \dot{x} + \partial_y X \, \dot{y} + \partial_t X \ , \tag{2}$$

$$V = \partial_x Y \, \dot{x} + \partial_y Y \, \dot{y} + \partial_t Y \ , \tag{3}$$

$$W = \partial_x Z \, \dot{x} + \partial_y Z \, \dot{y} + \partial_t Z \ . \tag{4}$$

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:

$$\frac{\partial(Z,Y)}{\partial(x,y)} U + \frac{\partial(X,Z)}{\partial(x,y)} V + \frac{\partial(Y,X)}{\partial(x,y)} W + \frac{\partial(X,Y,Z)}{\partial(x,y,t)} = 0 \ , \tag{5}$$

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 practice many sensors have aligned world and sensor coordinate systems which implies $\partial_y X = \partial_x Y = 0$. Yet Eq. (5) poses the general constraint independent of a particular sensor.

2.2 TLS Solution

Equation (5) poses only one constraint in the three unknowns U,V,W. This manifestation of the aperture problem has been examined in more detail before [7]. 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. (5) are bound to contain errors it is reasonable to use total least squares estimation as opposed to standard least squares. To do so

we rewrite Eq. (5) as $\boldsymbol{d}^T \boldsymbol{p} = 0$ with a data vector $\boldsymbol{d}$ given by the Jacobians. In order to avoid the trivial solution we require that $|\boldsymbol{p}| = 1$. It is straightforward to show that the solution is given by the eigenvector corresponding to the smallest eigenvalue of the so called structure tensor [7]:

$$J_{ij} = B * (d_i \cdot d_j) \ , \quad i,j = 1 \ldots 4 \ . \tag{6}$$

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:

$$\boldsymbol{f} = (U\,V\,W)^T = \frac{1}{p_4}(p_1\,p_2\,p_3)^T \ . \tag{7}$$

The smallest eigenvalue gives the residual of the attempted fit. This can be used to define a confidence measure based on a threshold τ [9]:

$$\omega = \begin{cases} 0 & \text{if } \lambda_4 > \tau \\ \left(\frac{\tau - \lambda_4}{\tau + \lambda_4}\right)^2 & \text{else} \end{cases} \ . \tag{8}$$

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, how this is done is described next.

2.3 Including Intensity

The usage of intensity data in addition to the range data can improve both the accuracy and density of the estimated range flow significantly [8]. 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:

$$0 = \partial_x I\,\dot{x} + \partial_y I\,\dot{y} + \partial_t I \ . \tag{9}$$

Combined with (2) and (3) we obtain an additional constraint on U and V:

$$\frac{\partial(I,Y)}{\partial(x,y)}U + \frac{\partial(X,I)}{\partial(x,y)}V + \frac{\partial(X,Y,I)}{\partial(x,y,t)} = 0 \ . \tag{10}$$

This can also be written as $\boldsymbol{d'}^T \boldsymbol{p} = 0$, where we simply set $d'_3 = 0$. The intensity constraint (10) results in another structure tensor $\boldsymbol{J'}$ constructed following Eq. (6). The sum of the two tensors yields a combined tensor from which the solution is then found by the analysis described above.

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 $\boldsymbol{J'}$ to account for different signal to noise ratios.

2.4 Dense Flow Fields

Even the usage of intensity data does not ensure a unique solution in every local neighbourhood. Because of the need to compute derivatives of both the range data and range flow they are required to vary smoothly. We employ normalized averaging to obtain the required smoothness. This averaging is a special case of normalized convolution and is computed as follows [3]:

$$\bar{\boldsymbol{f}} = \frac{B * (\Omega \cdot \boldsymbol{f})}{B * \Omega} \ . \tag{11}$$

Ω contains a confidence value for each estimated flow according to Eq. (8). The range data is smoothed in the same way where the certainty of the range measurement is used as confidence values.

3 Expansion Rates

There are a number of applications where one wants to determine whether the movement and deformation of an observed surface changes its surface area. Examples are the study of the effects of temperature changes or mechanical stress on materials and the observation of growth rates in biological systems.

The surface area of a regular surface can be defined as the integral over the area of tangential parallelograms [2]:

$$A = \int |\partial_x \boldsymbol{r}(x,y) \times \partial_y \boldsymbol{r}(x,y)| \ dxdy \ . \tag{12}$$

Hence we can define a local area element as $|\partial_x \boldsymbol{r} \times \partial_y \boldsymbol{r}|$. The relative change in the local surface area caused by the displacement vector $\Delta \boldsymbol{r} = \Delta t \boldsymbol{f}$ is then given by:

$$dA = \frac{|\partial_x(\boldsymbol{r} + \Delta \boldsymbol{r}) \times \partial_y(\boldsymbol{r} + \Delta \boldsymbol{r})|}{|\partial_x \boldsymbol{r} \times \partial_y \boldsymbol{r}|} \ . \tag{13}$$

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

$$e = (dA - 1) \cdot 100 \ \% \ . \tag{14}$$

This quantity can again be computed from derivatives, obtained via convolution, of both the range data and range flow arrays.

Because of the need to compute derivatives of both the range data and range flow they are required to vary smoothly and should be on the same scale. Thus we employ the same normalized averaging on the range data as on the range flow, see Sect. 2.4.

4 Experiments

In the following all derivatives are computed using 5-tap derivative filters that have been designed to give an optimal direction estimation [5]. For the integration in the structure tensor, see Eq. (6), we use a 9×9 Binomial. The threshold on the smallest eigenvalue is set to $\tau = 0.1$ and averaging in Eq. (11) is done by computing the second level of the Gaussian pyramid.

4.1 Synthetic Data

The synthetic data used here to test the introduced algorithm consists of an expanding sphere. The sensor is taken to use a pinhole camera, as is the case for structured lighting or stereo systems. The radius of the sphere is initially 150mm and its centre is located at $(0, 0, 300)^T\, mm$. The focal length is set to 20mm and sensor elements (CCD pixel) are placed at 0.05mm in both directions. We apply an intensity texture to each surface element based on the spherical angles $I = I(\theta, \phi)$:

$$I = \begin{cases} 100 & \text{if } \theta < 0.5^\circ \\ 100 + 50\sin(\frac{2\pi\theta}{1^\circ}) + 50\sin(\frac{2\pi\phi}{30^\circ}) & \text{else} \end{cases} . \tag{15}$$

Surface expansion is modelled by a change in radius. We use a sequence with an expansion rate of 1% per frame, this corresponds to multiplying the radius with a factor of 1.00499 between successive frames. Additionally the whole sphere is moving with a constant velocity of $(0.01, 0.02, 0.03)^T\, [mm/frame]$.

Sensible values for the noise in the range data with a viewing distance around 150mm are given by $\sigma_Z = 0.1mm$ [1]. For the error in X and Y coordinates we assume a noise around $\sigma_{X,Y} = 0.01mm$ and the intensity noise lies in the range of $\sigma_I = 1.0$.

The following table gives the mean values for the relative error in the expansion rate E_e, the relative error in the range flow magnitude E_m and the mean absolute angle deviation between correct and estimated 3D velocity E_d:

$\sigma_{X,Y}$	σ_Z	σ_I	E_e [%]	E_m [%]	E_d [°]
0.00	0.0	0.0	1.02	0.001	0.01
0.01	0.0	0.0	1.35	0.002	0.02
0.00	0.1	0.0	1.10	0.001	0.01
0.00	0.0	1.0	3.24	0.003	0.05
0.02	0.0	0.0	2.03	0.003	0.04
0.00	0.2	0.0	1.55	0.001	0.02
0.00	0.0	2.0	6.32	0.004	0.10
0.01	0.1	1.0	3.11	0.003	0.05
0.02	0.2	2.0	6.89	0.005	0.10

First we conclude that the range flow field is very accurate even for slightly higher noise levels. However the estimation of expansion rates depends strongly

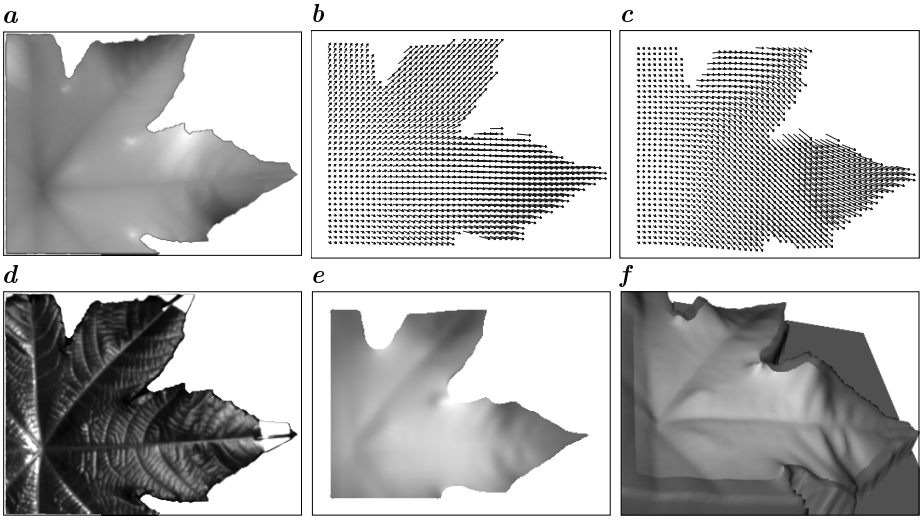

Fig. 1. Real castor bean leaf: **a** depth (Z) data, **b** U,V-component and **c** U,W-component of the range flow field. **d** Intensity data. Growth in the range from -2% to 3% per hour: **e** as map and **f** as texture on the 3D structure.

on the noise. For a standard deviation of 0.1mm in the Z coordinate we can compute accurate expansion rates. This accuracy can be achieved with current sensor technology.

4.2 Leaf Growth

An application that requires the evaluation of expansion rates is the study of growth in leaves. As an example we investigate a moving and growing castor bean leaf. The 3D structure is captured using a structured light range sensor. We sample the leaf every 2 minutes to obtain a sequence of range data sets. Figure 1a,d show the depth and intensity data for the frame where we compute range flow. The flow field is given in Fig. 1b,c. We see that there is considerable motion of the leaf in particular in Z direction. Clearly a lot of this motion does not stem from growth alone. The obtained growth rate in % per hour are given in Fig. 1e,f. In accordance with previous findings from leaves that have been confined to a plane we see that growth diminishes as we move from the base to the tip of the leaf [6].

5 Conclusion

We introduced a new version of the range flow motion constraint equation that can be evaluated directly on the sensor grid. In particular there is no need to fit a surface model to the range data. A formula for the computation of expansion rates from the thus computed range flow fields has been given. We could demonstrate that our method is capable to compute accurate range flow and surface expansion rates on both synthetic and real data.

Bibliography

[1] 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, Austria, October 1995.

[2] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976.

[3] G. H. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic, Dordrecht, The Netherlands, 1995.

[4] B. K. P. Horn and J. Harris. Rigid body motion from range image sequences. *CVGIP*, 53(1):1–13, January 1991.

[5] B. Jähne, H. Scharr, and S. Körkel. Principles of filter design. In *Handbook of Computer Vision and Applications*. Academic Press, 1999.

[6] D. Schmundt, M. Stitt, B. Jähne, and U. Schurr. Quantitative analysis of local growth rates of dicot leaves at high temporal and spatial resolution, using image sequence analysis. *Plant Journal*, 16:505–514, 1998.

[7] H. Spies, H. Haußecker, B. Jähne, and J. L. Barron. Differential range flow estimation. In *DAGM*, pages 309–316, Bonn, Germany, September 1999.

[8] 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.

[9] H. Spies, B. Jähne, and J. L. Barron. Regularised range flow. In *ECCV*, pages 785–799, Dublin, Ireland, June/July 2000.

[10] L. Tsap, D. Goldgof, and S. Sarkar. Multiscale combination of physically-based registration and deformation modeling. In *CVPR*, pages 422–429, June 2000.

[11] S. Vedula, S. Baker, P. Rander, R. Collins, and T. Kanade. Three-dimensional scene flow. In *ICCV*, pages 722–729, Pittsbrugh, PA, September 2000.

[12] 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.