Range Flow: The 3D Movement of Deformable Surfaces

Hagen Spies

IWR
University of Heidelberg
Heidelberg, Germany
http://klimt.iwr.uni-heidelberg.de/~hspies
Hagen.Spies@iwr.uni-heidelberg.de

John Barron

Department of Computer Science
University of Western Ontario
London, Ontario, Canada
http://www.csd.uwo.ca/faculty/barron
barron@csd.uwo.ca

This contribution describes how the instantaneous 3D-velocity field of a moving surface can be recovered from sequences of depth maps or range data sets. This is essentially an extension of optical flow to 3D-surface data and has been termed range flow.

1 Overview

We can compute local three-dimensional velocity (range flow) of moving surfaces from sequences of depth maps. The corresponding constraint equation (Sect. 2) is closely related to the optical flow constraint equation. Thus similar schemes can be employed to solve for the velocity field. However, in three dimensions we can distinguish three flavours of the aperture problem (Sect. 3). Sometimes, including intensity as an aditional data channel can help to overcome the aperture problem.

As example solution strategies, the use of a local total least squares estimation (Sect. 4) and a global regularisation framework (Sect. 5) are discussed. As an example application, we show how the motion of leaves can be captured (Sect. 7).

2 Range Flow Constraint Equation

First we note that the observed surface can be described by it's depth as a function of space and time $Z = Z(X, Y, t)$. The total derivative with respect to time directly yields the range flow motion constraint equation:

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

Here $\boldsymbol{f} = [U, V, W]^T$ is the range flow and indices denote partial derivatives. The use of only the first order expresses the assumption that the object under consideration can approximated by local planar patches. If needed, this can easily be extended to incorporate higher terms. The range flow motion constraint equation is the analogon to the *brightness change constraint* equation used for optical flow.

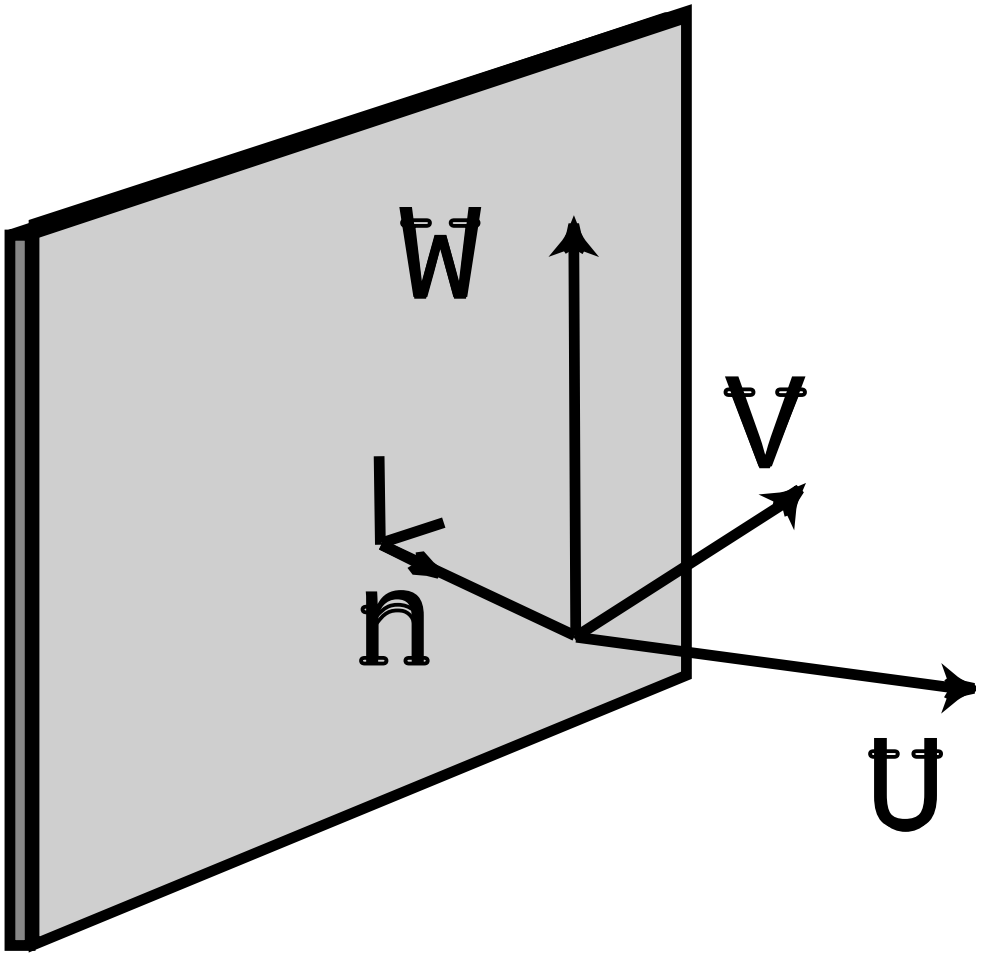
w
v
n
u

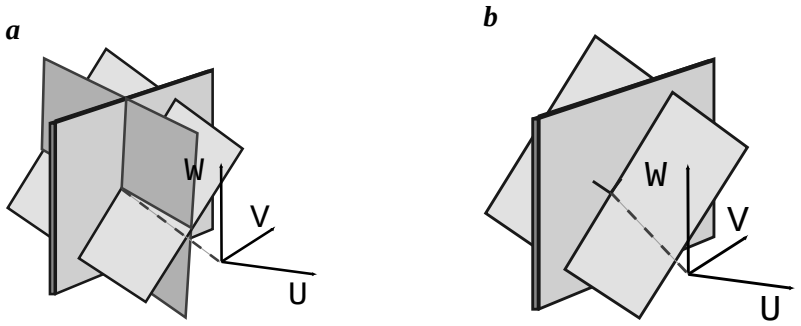

Figure 2: Illustration of the constraint equations present in a neighbourhood for **a** full flow and **b** line flow.

be the minimal vector between the origin and the constraint plane, see Fig. 1 . This gives the *raw normal flow*:

$$\boldsymbol{f}_r = \frac{-Z_t}{Z_X^2 + Z_Y^2 + 1} \begin{bmatrix} Z_X \\ Z_Y \\ 1 \end{bmatrix} . \tag{5}$$

As we are dealing with three dimensional flow, at least three mutually distinct, i.e. non-parallel, planes are needed to combine the various normal flows in a region to yield full 3D-flow. If there are three or more linearly independent constraint equations in the considered neighbourhood the full velocity vector is readily computed by the intersection of the constraint planes as illustrated in Fig. 2a. Section 4 describes how this can be done in a local neighbourhood by means of a total least squares technique.

Linear structures such as intersecting planes result in two constraint planes, see Fig. 2b. The point on the common line closest to the origin gives an appropriate normal flow called line flow. This line flow lies in the plane perpendicular to the linear structure. In this case only the motion along the direction of the intersecting line can not be resolved.

We now proceed to illustrate two example methods, local and global, that can be used for a full range flow estimation.

4 Local TLS Range Flow

The TLS solution presented here is an extension of the *structure tensor* algorithm for optical flow estimation.

Assuming constant flow in a region containing n pixel we have n equations (1). With $\vec{d} = [Z_X\ Z_Y\ 1\ Z_T]^T$, $\vec{u} = [U\ V\ W\ 1]^T$ and the data matrix $\boldsymbol{D} = [\vec{d}_1 \dots \vec{d}_n]^T$, the flow estimation in a total least squares sense can be formulated as:

$$||\boldsymbol{D}\vec{u}||_2 \rightarrow \min \quad \text{subject to} \quad \vec{u}^T\vec{u} = 1 \ . \tag{6}$$

The solution is given by the eigenvector $\hat{e}_4$, corresponding to the smallest eigenvalue λ_4 of the generalised structure tensor:

$$\boldsymbol{J} = \boldsymbol{D}^T\boldsymbol{D} = \begin{bmatrix} < Z_X Z_X > & < Z_X Z_Y > & < Z_X > & < Z_X Z_T > \\ < Z_Y Z_X > & < Z_Y Z_Y > & < Z_Y > & < Z_Y Z_T > \\ < Z_X > & < Z_Y > & < 1 > & < Z_T > \\ < Z_T Z_X > & < Z_T Z_Y > & < Z_T > & < Z_T Z_T > \end{bmatrix} . \quad (7)$$

Here $< \cdot >$ denotes local averaging using a Box or Binomial filter. The desired range flow is then given by:

$$\vec{f}_f = \frac{1}{e_{44}} \begin{bmatrix} e_{14} \\ e_{24} \\ e_{34} \end{bmatrix} . \quad (8)$$

As $\boldsymbol{F}$ is real and symmetric the eigenvalues and eigenvectors can easily be computed using Jacobi-Rotations. In order to save execution time we only compute range flow where the trace of the tensor exceeds a threshold τ_1. This eliminates regions with insufficient magnitude of the gradient.

In the above we are really fitting a local constant flow model to the data. The smallest eigenvalue λ_4 directly measures the quality of this fit.

4.1 Normal Flows

Let the eigenvalues of $\boldsymbol{F}$ be sorted: $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \lambda_4$. Thus if $\lambda_3 \approx \lambda_4$ no unique solution can be found. In general, any vector in the nullspace of $\boldsymbol{F}$ is a possible solution. In this case it is desirable to use the solution with minimal norm.

On planar structures all equations (1) are essentially the same. Only the largest eigenvalue is significantly different ($> \tau_2$) from zero. The so called plane flow can then be found from the corresponding eigenvector $\hat{e}_1 = [e_{11}\ e_{21}\ e_{31}\ e_{41}]^T$ as follows:

$$\boldsymbol{f}_p = \frac{e_{41}}{e_{11}^2 + e_{21}^2 + e_{31}^2} \begin{bmatrix} e_{11} \\ e_{21} \\ e_{31} \end{bmatrix} . \quad (9)$$

Linear structures exhibit two types of constraints within the considered aperture, the minimum norm solution (line flow) is found from the eigenvectors $\hat{e}_1,\ \hat{e}_2$:

$$\boldsymbol{f}_l = \frac{1}{1 - e_{41}^2 - e_{42}^2} \left[e_{41} \begin{bmatrix} e_{11} \\ e_{21} \\ e_{31} \end{bmatrix} + e_{42} \begin{bmatrix} e_{12} \\ e_{22} \\ e_{32} \end{bmatrix} \right] . \quad (10)$$

Figure 3 shows an example of the various flow types.

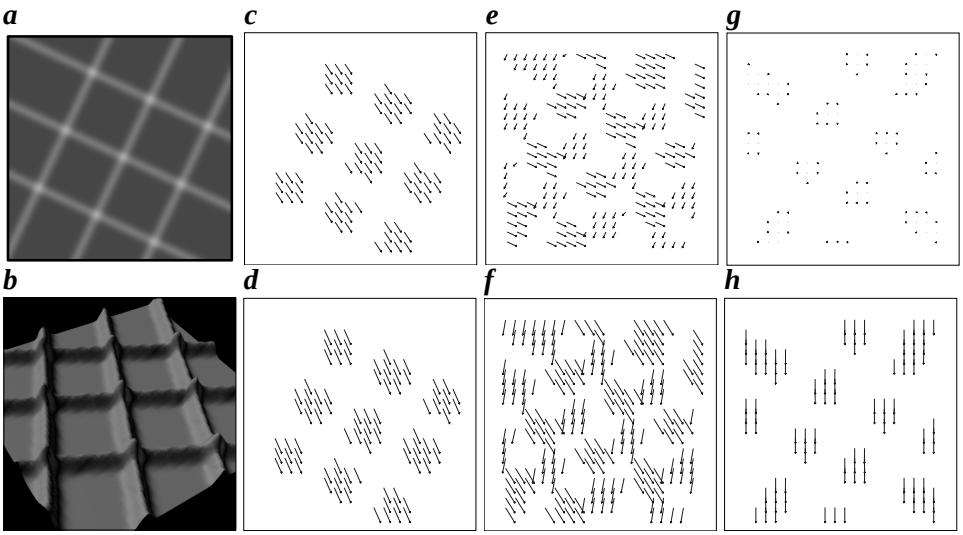

Figure 3: Example flow types: **a** synthetic depth map, **b** rendered. $X-Y$ components of the estimated flow fields: **c** full flow, **e** line flow and **g** plane flow and $X-Z$ components of the estimated flow fields: **d** full flow, **f** line flow and **h** plane flow.

5 Globally Smooth Range Flow

Another way of solving the range flow estimation problem is to use a variational approach. In analogy to the well known optical flow algorithm by Horn and Schunk we can impose a smoothness constraint on the flow field. To ensure smoothly varying parameters we use as the smoothness term:

$$\sum_{i=1}^{3}(\nabla \boldsymbol{f}_i)^2 \rightarrow \min \ . \tag{11}$$

Obviously the use of this simple membrane model is only justified if we have already segmented the data into differently moving objects. If no such segmentation were available more elaborate schemes would have to be considered. The minimisation combining data (1) and smoothness (11) terms reads:

$$\int_A \left\{ (Z_X U + Z_Y V + W + Z_t)^2 + \alpha^2 \sum_{i=1}^{3}(\nabla \boldsymbol{f}_i)^2 \right\} \mathrm{d}r \rightarrow \min . \tag{12}$$

Here α can be used to regulate the influence of the smoothness term. Just as in the optical flow case this can be solved iteratively via the Euler-Lagrange equations. Figure 4 shows the result for increasing iteration time on the sequence shown in Fig. 3. While the flow is initially dominated by the aperture problem it increasingly tends towards a smooth flow field. However, depending on the data, convergence can be rather slow.

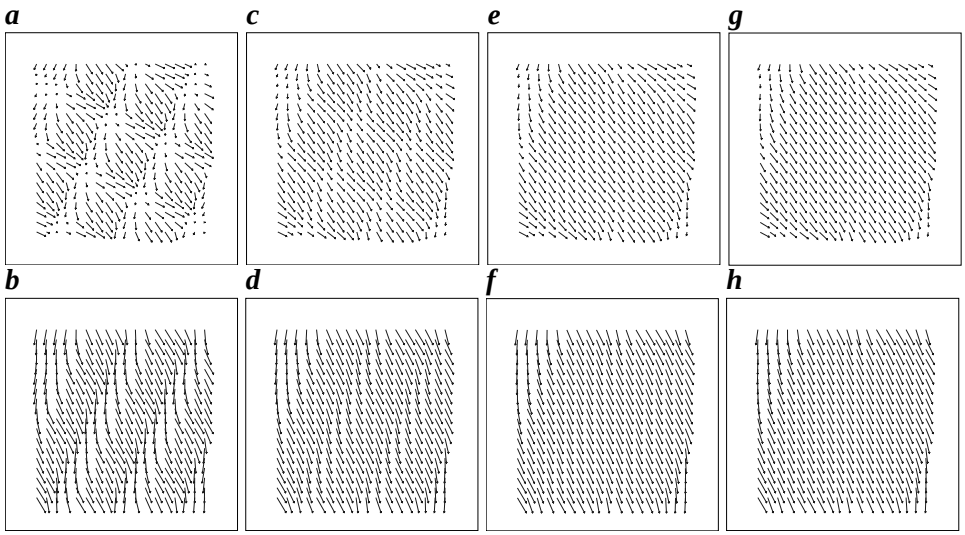

Figure 4: Global smoothness result for different number of iterations: $X-Y$ components of the estimated flow fields: **a** 10, **c** 50, **e** 100 and **g** 1000 iterations. $X-Z$ components of the estimated flow fields: **b** 10, **d** 50, **f** 100 and **h** 1000 iterations.

6 Intensity as Additional Data Channel

When the range data is obtained from stereo, structured lighting, depth from X or active laser triangulation we typically have some intensity data as well. As this intensity is perfectly registered with the depth data it can be used to aid the range flow estimation.

Assuming the intensity does not change with moderate depth changes we obtain an additional constraint equation in the X and Y movements:

$$I_X U + I_Y V + I_t = 0 \ . \tag{13}$$

Together with the range flow motion constraint (1) we now have two equations in three unknowns. Even if they are independent, which is not always the case, we still need to make further assumptions to solve for full range flow.

For the local TLS estimation we note that Eq. (13) can also be written in the form of $\boldsymbol{d'}^T \boldsymbol{u} = 0$, where we simply set $d'_3 = 0$. Thus the intensity constraint results in another structure tensor $\boldsymbol{J'}$ constructed following Eq. (7). The sum of the two tensors yields a combined tensor from which the solution is then found by the analysis described above, see Sect. 4.

In the case of the global smoothness algorithm (Sect. 5) we simply include the intensity constraint in the energy functional:

$$\int_A \left\{ (Z_X U + Z_Y V + W + Z_t)^2 + \beta (I_X U + I_Y V + I_t)^2 + \alpha^2 \sum_{i=1}^{3} (\nabla \boldsymbol{f}_i)^2 \right\} \mathrm{d}r \rightarrow \min \ . \tag{14}$$

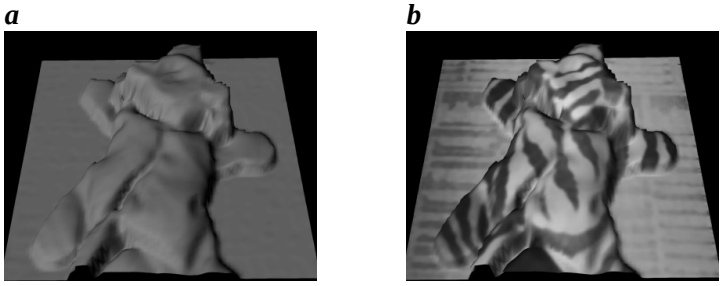

Figure 5: Tiger test sequence: a) depth only and b) with intensity texture.

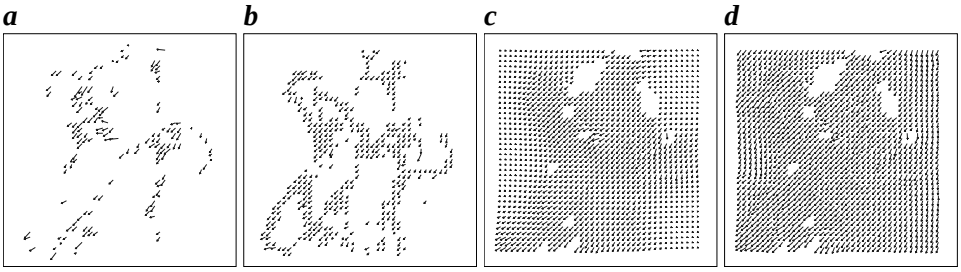

Figure 6: X-Y-range flow. TLS: a) depth only, b) with intensity. Global Smoothness (500 iterations): c) depth only, d) with intensity.

Here we introduced a relative weight β that can be used to give more influence to either the depth or intensity data. A solution can again be found from the Euler-Lagrange equations.

An example might illustrate the performance gain that can be achieved by incorporating the intensity information. Figure 5 shows data taken with a laser range finder from a toy tiger that is moved on a system of linear positioners. Due to rapid depth changes there are some holes in the data, these are left out in the computation.

The X-Y components of the computed full flow fields are shown in Fig. 6. The increased density in the TLS case is apparent. Some numbers may illustrate the algorithms behavior:

data	density [%]	rel mag error [%]	dir error [°]
Z only	10.5	13.8 ± 15.4	12.7 ± 10.3
I and Z	59.0	7.9 ± 9.0	9.9 ± 8.6

Not only does the density increase but also the error decreases.

Figure 6c,d indicate that the combined use of intensity and depth in the global flow algorithm also achieves a better estimate after the same number of iterations.

7 Application Example: Plant Leaf Motion

We finally present some flow fields found by observing living castor oil leaves. Figure 7 shows four examples of the type of data and flow fields encountered in

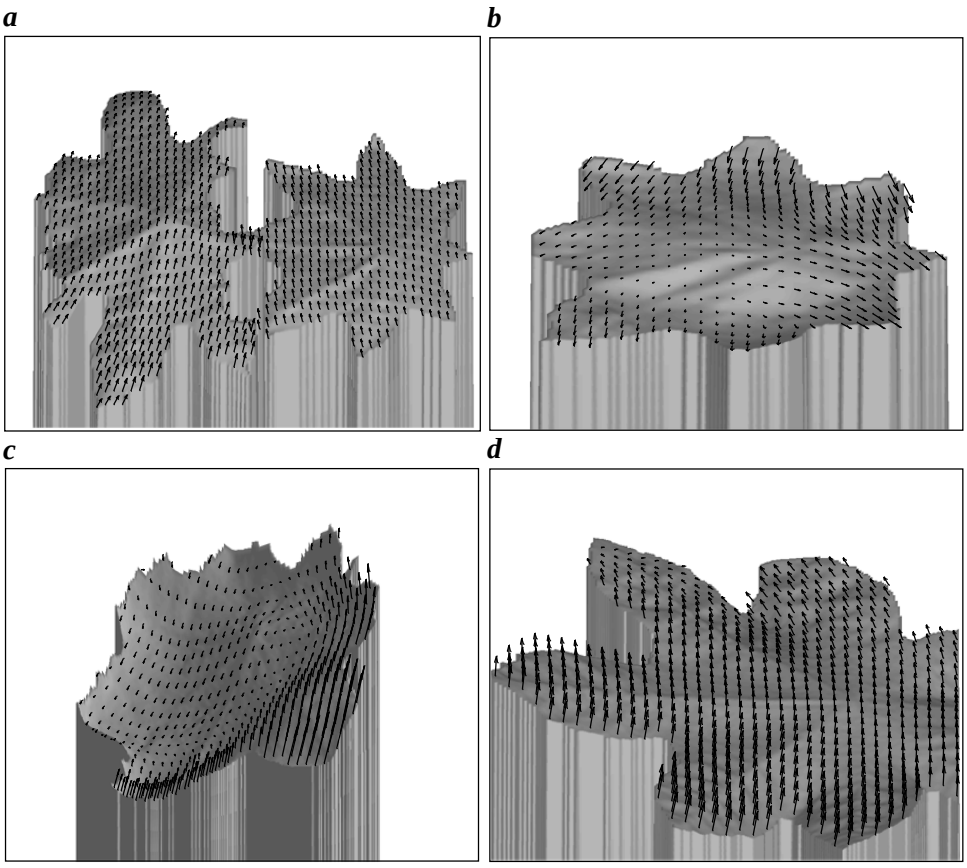

Figure 7: Example movements of castor oil plant leaves.

this application. The folding of the outer lobes is clearly visible as well as a fair bit of lateral motion of the leaf. The data sets considered here are taken at night with a sampling rate of 5 minutes.

In collaboration with the botanical institute at the University of Heidelberg and the Agriculture and Agri-Food Canada research station in Harrow, Ontario we seek to establish the diurnal motion patterns of such leaves.

8 More Information on Range Flow

The interested reader can find more a detailed description of this work in some papers at:

- http://klimt.iwr.uni-heidelberg.de/ hspies/publications.html
- http://www.csd.uwo.ca/faculty/barron