Principal Components Analysis of Optical Snow

V. C.-Couture[1], S. Roy[2], M. S. Langer[3] and R. Mann[4]

[1] University of Montreal (DIRO), Montreal, Quebec, H3T 1J4 Canada

email: chapdelv@iro.umontreal.ca

[2] University of Montreal (DIRO), Montreal, Quebec, H3T 1J4 Canada

[3] McGill University (School of Computer Science), Montreal, Quebec, H3A 2A7 Canada

[4] University of Waterloo (School of Computer Science), Waterloo, Ontario, N2L 3G1 Canada

Abstract

Many applications in computer vision use Principal Components Analysis (PCA), for example, in camera calibration, stereo, localization and motion estimation. We present a new and fast PCA-based method to analyze optical snow. Optical snow is a complex form of visual motion that occurs when an observer moves through a highly cluttered 3D scene. For this category of motion field, no spatial or depth coherence can be assumed. Previous methods for measuring optical snow have used a wedge filter in a spatiotemporal frequency domain. The PCA method is also based on the spatiotemporal frequency domain analysis, but examines a different geometry property of the spectrum. We compare the results of the PCA method to the previous methods using both real and synthetic sequences.

1 Introduction

Recently, [12] considered a category of visual motion called *optical snow* which generalizes optical flow by abandoning classical assumptions of spatial continuity. Optical snow arises when an observer moves relative to a rigid 3-D cluttered scene or object (falling snow, forest, plants). Optical snow is characterized by small spatial features and a dense set of depth discontinuities. It produces a highly discontinuous motion field. Traditional optical flow methods cannot be expected to recover an image velocity field from optical snow since these methods typically assume local smoothness in the velocity field [1], or a small number of well-isolated discontinuities e.g. [19, 20, 3]. By definition, smoothness constraints do not apply for optical snow. Layered motion models such as [2, 10, 11, 5, 4] also do not apply, since these models assume smoothness within layers[24], and the number of layers is typically very small (2 or 3). In optical snow, the number of "layers" can be in the hundreds or more.

To overcome the dense depth discontinuity problem that arises in measuring optical snow, [12] introduced an analysis of the motion which was based in the spatio-temporal frequency domain. [13, 14] presented an algorithm to find the direction of the motion for the special case of parallel optical snow, which arises for example in the case of a lateral camera motion [13, 15]. This algorithm was then generalized in [18] to handle the case of optical snow that arises from general observer motion, including both non-lateral motion and camera roll.

The previous method for measuring optical snow uses a wedge-filter to estimate the distribution of power in the spatio-temporal Fourier power spectrum. This filter was motivated by the geometric properties of optical snow in the spatio-temporal frequency domain (see Sec. 2). In this paper, we present a new and fast algorithm which uses PCA to estimate the properties of optical snow. The PCA method is complementary to the wedge filter method, in that it is based on a geometrical property of the power spectrum that the wedge method ignores. We show the effectiveness of the PCA method by comparing experimental results of the PCA method to that of the wedge method of [14].

2 Previous work

2.1 Optical Snow

Optical snow is produced by an observer moving relative to a rigid 3D cluttered scene. It is a special case of *motion parallax* in which the density of depth discontinuities is high. For motion parallax in general (and optical snow in particular), the image velocity field can be closely approximated as a sum of two fields – one that is due to camera translation and the other due to camera rotation [16]. As a result, the image velocity vectors in a local image patch satisfy the relation

$$(v_x, v_y) = (u_x + \alpha\, t_x,\ u_y + \alpha\, t_y) \tag{1}$$

where u_x, u_y, t_x, t_y are constants which depend on the patch and α depends on position x,y in the patch and on the depth of the point visible at (x,y) [18]. Thus, optical snow yields a one parameter family of velocities (a line!) in any local image patch.

One can examine the power spectrum properties of optical flow by using Eq. (1) to extend the classical motion plane property [23, 7]. The motion plane property states that an image translating with uniform image velocity produces a plane of energy in the 3D spatio-temporal frequency domain. Formally, let $I(x,y,t)$ be a time varying image which is defined by an image patch translating with velocity (v_x, v_y). From [8], this velocity is constrained by

$$v_x \frac{\partial I}{\partial x} + v_y \frac{\partial I}{\partial y} + \frac{\partial I}{\partial t} = 0\ \ . \tag{2}$$

This constraint, transposed in the Fourier domain, yields

$$(v_x f_x + v_y f_y + f_t)\ \ \hat{I}(f_x, f_y, f_t)\ =\ 0 \tag{3}$$

where $\hat{I}(f_x, f_y, f_t)$ is the Fourier transform of $I(x,y,t)$. Eq. (3) implies that all frequencies (f_x, f_y, f_t) for which $\hat{I}(f_x, f_y, f_t) \neq 0$ lie on the plane

$$v_x f_x + v_y f_y + f_t = 0\ . \tag{4}$$

Substituting Eq. (1) into Eq. (4) yields a family of planes in the frequency domain,

$$(u_x + \alpha\, t_x) f_x\ +\ (u_y + \alpha\, t_y\ f_y + f_t)\ =\ 0$$

This set of planes forms a *bowtie* (see Fig. 1a) and follows the two following propositions [14]:

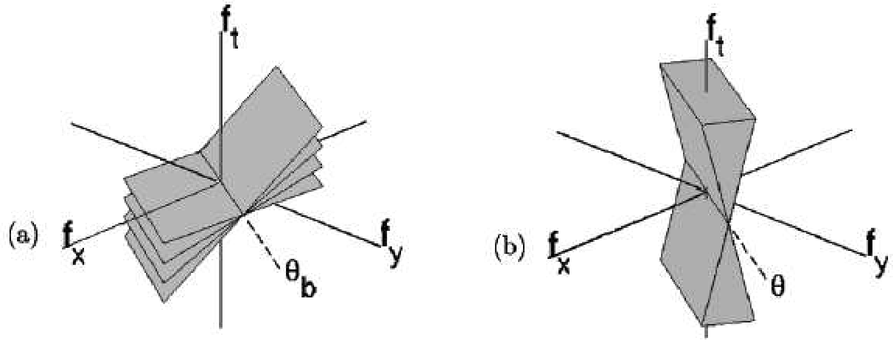

Figure 1: (a) Bowtie signature in the frequency domain. The axis of bowtie is in direction θ_b. (b) Wedge used to estimate the orientation of the bowtie. The dotted line at angle θ in the (f_x, f_y) is denoted l_θ.

Proposition 1: The planes of the bowtie intersect at a common line, called the *axis of the bowtie*, that passes through the origin.
Proposition 2: The axis of the bowtie is in direction $(-t_y, t_x, u_x t_y - u_y t_x)$ in the 3D frequency domain. By normalizing (t_x, t_y), this direction becomes $(-t_y, t_x, |U| sin(\phi))$, where ϕ is the angle between vectors (t_x, t_y) and (u_x, u_y). By assuming that (t_x, t_y) is perpendicular to (u_x, u_y), we get $(-t_y,\ t_x,\ \sqrt{u_x^2 + u_y^2}\ )$.

For simplicity, all equations in this paper assume image sequences of equal dimension in space and time, i.e. the size N of image region equals the number T of frames. In practice, the case of unequal dimensions can be easily be accounted for using factors of $\frac{N}{T}$.

We emphasize that optical snow is a very general model of visual motion. It describes the velocities of *any* rigid 3D scene as seen by a moving observer [14, 18]. The model applies whether the surfaces are partly transparent, smooth, layered, or densely cluttered. The model does not assume any spatial continuity. This is in sharp contrast to conventional optical flow estimation which relies on spatial coherence of the motion fields or uses motion layers model.

2.2 Bowtie axis estimation using wedge filter

[12, 14] introduced an algorithm to find the bowtie axis in the case that $(u_x, u_y) = (0, 0)$, where no rotation is allowed. In this case, the bowtie axis lies in the $(f_x, f_y, 0)$ plane and the direction of the bowtie axis is $(-t_y, t_x, 0)$. To find the bowtie axis, a *wedge filter* is used. The wedge filter is defined by two motion planes, of slope $\pm v_{max}$, respectively. These two planes intersect at a line in the (f_x, f_y) plane, which is oriented at an an angle θ (see figure 1). We let l_θ denote this line. It corresponds exactly to the axis of the bowtie when $\theta = \theta_b$. If all motion planes in the bowtie have slopes of magnitude that is less than v_{max} then, for this θ, the intersection of the wedge filter and the bowtie will be restricted entirely to the bowtie axis.

To find the bowtie axis, the wedge filter is rotated (θ is varied) and the power that falls within the wedge filter is measured as a function of θ. In particular – and this is an important detail – a small cylinder containing the line l_θ is removed from the wedge filter. This remove all power from the wedge filter for the case that $\theta = \theta_b$. The angle θ_b of the bowtie axis is estimated to be the angle for which the power in the wedge filter is a minimum.

The more general case is that $(u_x, u_y) \neq (0,0)$. In this case, the bowtie axis does not lie in the (f_x, f_y) plane (see Proposition 2). To estimate the bowtie axis in this case, a method was presented in [18]. This method has two steps. The first step is to find the best fit motion plane for the bowtie and then shear the power spectrum in the f_t direction to bring this motion plane in correspondence with the (f_x, f_y) plane. This shear amounts to *motion compensation* [17] of the image sequence (within the patch in question) such that the mean velocity in the patch becomes 0. In particular, since the bowtie axis is contained in all motion planes within the bowtie, this best fit plane also contains the bowtie axis. The second step of the method for finding the bowtie is to find the θ that best aligns the wedge filter to the bowtie, as described above for the case that the bowtie axis happens to already lie in the (f_x, f_y) plane.

To perform motion compensation, [18] finds the vector (v_x, v_y) that minimizes the sum of squares:

$$g(v_x, v_y) \;=\; \sum_{f_x, f_y, f_t} \mid \hat{I}(f_x, f_y, f_t) \mid^2 \; (f_t - (v_x f_x + v_y f_y) \; mod \; N))^2$$

We re-express the motion compensation step, in a way that will allow us to compare it to the PCA method, which we introduce in this paper.

Let $\mathbf{F}$ be the $N^3 \times 3$ matrix which lists the N^3 possible frequency triplets (f_x, f_y, f_t), where each f_x, f_y, f_t value is in $\{0, \ldots, N-1\}$. Let $\mathbf{W}$ be an $N^3 \times N^3$ diagonal matrix of weights, with diagonal elements,

$$W(f_x, f_y, f_t) \;=\; |\hat{I}(f_x, f_y, f_t)| \;.$$

Consider the matrix product

$$\mathbf{F}^T \mathbf{W}^2 \mathbf{F} \;=\; \begin{pmatrix} \Sigma f_x^2 \, W(f_x, f_y, f_t)^2 & \Sigma f_x \; f_y \, W(f_x, f_y, f_t)^2 & \Sigma f_x \; f_t \, W(f_x, f_y, f_t)^2 \\ \Sigma f_x f_y \, W(f_x, f_y, f_t)^2 & \Sigma f_y^2 \, W(f_x, f_y, f_t)^2 & \Sigma f_y \; f_t \, W(f_x, f_y, f_t)^2 \\ \Sigma f_x f_t \, W(f_x, f_y, f_t)^2 & \Sigma f_t f_y \, W(f_x, f_y, f_t)^2 & \Sigma f_t^2 \, W(f_x, f_y, f_t)^2 \end{pmatrix}$$

where each of the summations is a triple summation over f_x, f_y, f_t, namely over the entire 3D frequency domain. In the case that there is no temporal aliasing, the modulus N can be ignored and the summation is equivalent to

$$g(v_x, v_y) \;\equiv\; (v_x, v_y, -1) \, \mathbf{F}^T \mathbf{W}^2 \, \mathbf{F} \, (v_x, v_y, -1)^T$$

This method thus finds the best fitting motion plane by minimizing the weighted squared distance in the f_t direction.

With this background, let us now turn to the original contribution of this paper which is to introduce a new, simpler and faster PCA-method for finding the bowtie axis.

3 Bowtie Axis Estimation with PCA

The basic idea of our PCA method is to find a line in the frequency domain that passes through the origin and contains a concentration of power. It does so by computing the eigenvectors of a 3×3 covariance matrix which is defined in the (f_x, f_y, f_t) domain. This strategy is justified by the fact that the energy on the bowtie axis is the superposition of the energies (not merely the intersection) of the planes that contribute to the bowtie (see Fig. 2).

The PCA method uses classical *orthogonal distance regression* [21], which finds the best fit line and best fit plane to a set of 3D points. In our problem, the 3D space is the frequency domain and so the coordinates are (f_x, f_y, f_t). The lines and planes we are interested in pass through the origin. Let

$$n_x \, f_x \; + \; n_y \, f_y \; + \; n_t \, f_t \; = \; 0$$

be the equation of a plane that passes through the origin, where

$$\mathbf{n}^T = (n_x, n_y, n_t)$$

is a vector normal to the plane. We wish to minimize the *sum of squared orthogonal weighted distances* of each (f_x, f_y, f_t) to this plane,

$$f(\mathbf{n}) \; = \; \frac{1}{\mathbf{n}^T \mathbf{n}} \sum_{f_x, f_y, f_t} W(f_x, f_y, f_t)^2 \; |n_x \, f_x \; + \; n_y \, f_y \; + \; n_t \, f_t|^2 \; .$$

The function $f(\mathbf{n})$ can be expressed as the matrix product

$$f(\mathbf{n}) = (\mathbf{n}^T \; \mathbf{F}^T \; \mathbf{W}^2 \; \mathbf{F} \; \mathbf{n}) / (\mathbf{n}^T \mathbf{n})$$

where $\mathbf{F}$ is the same matrix as in Sec. 2. We will discuss our choice of $W(f_x, f_y, f_t)$ below.

The function $f(\mathbf{n})$ is a *Rayleigh quotient*. This quotient is maximized and minimized by the eigenvectors corresponding to the largest smallest eigenvalues of $\mathbf{F^T W^2 F}$, respectively. The minimum eigenvector defines the normal of the best fit plane π. The maximum eigenvector defines to the best fit line l which is considered to be the bowtie axis. The line l always lies in the plane π, since $\mathbf{F}^T\mathbf{W^2 F}$ is a symmetric real 3×3 matrix and so it has an orthogonal set of eigenvectors.

3.1 Comparison of the two best fit plane methods

Both the wedge filter method of [14, 18] and our new PCA method compute a minimization which finds a motion plane that best fits a bowtie. However, there are several differences between the methods.

Let us first briefly mention two minor differences. One is that $f(\mathbf{n})$ measures the weighted distance orthogonal to a plane, whereas $g(v_x, v_y)$ measures the weighted distance in the f_t direction only. One can make arguments for either of these distance choices. Indeed similar non-orthogonal [7, 22] vs. orthogonal [6, 9] choices were made in classical optical flow. A second difference is that the $f(\mathbf{n})$ minimization normalizes by the length of $\mathbf{n}$, whereas the $g(v_x, v_y)$ minimization does not normalize by the length of $(v_x, v_y, -1)$. The effect of normalizing is to explicitly penalize for high speeds. Again, arguments can be made for or against.

A more significant difference between the two methods is the choice of $W(f_x, f_y, f_t)$. For wedge filter method [18], the weighting function was the amplitude spectrum $|\hat{I}(f_x, f_y, f_t)|$. The amplitude spectrum turns out to be inappropriate for the PCA formulation, however.

To understand how problems can arise, consider an illustrative example. Suppose that an image sequence is defined by summing two translating image sequences

$$I(x,y,t) \ = \ I_1(x,y,t) + I_2(x,y,t) \ .$$

The Fourier transform of the summed image sequence is

$$\hat{I}(f_x, f_y, f_t) = \hat{I}_1(f_x, f_y, f_t) + \hat{I}_2(f_x, f_y, f_t)$$

with power spectrum

$$|\hat{I}(f_x, f_y, f_t)|^2 \ = \ |\hat{I}_1(f_x, f_y, f_t))|^2 + |\hat{I}_2(f_x, f_y, f_t))|^2 + 2\, Re(\hat{I}_1(f_x, f_y, f_t)\, \overline{\hat{I}_2(f_x, f_y, f_t)}\,).$$

Let's consider the expected value. The phases of $\hat{I}_1(f_x, f_y, f_t)$ and $\hat{I}_2(f_x, f_y, f_t)$ can be assumed independent, so we have

$$\mathscr{E}\ \{\ \hat{I}_1(f_x, f_y, f_t)\, \overline{\hat{I}_2(f_x, f_y, f_t)}\ \}\ = 0$$

and so

$$\mathscr{E}\ \{\ |\hat{I}(f_x, f_y, f_t)|^2\ \}\ = \mathscr{E}\ \{\ |\hat{I}_1(f_x, f_y, f_t))|^2\ \} + \mathscr{E}\ \{\ |\hat{I}_2(f_x, f_y, f_t))|^2\ \}\ .$$

The expected power spectra on the bowtie axis is thus the sum of the expected power spectrum of the two component image sequences. Thus there is a concentration of power on the bowtie axis, which is where the two planes superimpose.

Unfortunately, there is no reason why the PCA method should detect this concentration of power, if we choose $W(f_x, f_y, f_t) = |\hat{I}(f_x, f_y, f_t)|$ as does the wedge method. To illustrate, take the counter-example in which the first frame (and hence all frames) of the two translating images have the same amplitude spectra, $|\hat{I}_1(f_x, f_y)| \ = \ |\hat{I}_2(f_x, f_y)|$, but they have independent phase spectra. Further suppose that two translating images sequences have equal but opposite velocities. Then one can show that two of the eigenvectors of $\mathbf{F}^T\mathbf{W}^2\mathbf{F}$ lie in the plane (f_x, f_y) and the third eigenvector is parallel to f_t. Moreover, the direction of the two eigenvectors in (f_x, f_y) depends on only on the amplitude spectra, $|\hat{I}_1(f_x, f_y)| \ = \ |\hat{I}_2(f_x, f_y)|$. These eigenvector directions would be independent of the direction of image motion. Thus, for this counterexample, the PCA method would *not* be able to recover the bowtie axis, even though there *is* a concentration of power there. We have done experiments to show that this difficulty arises for more general sequences. (This data are not included because of space limitations.)

3.2 Finding the bowtie axis

The difficulty that we just described arises because the power spectrum is a second order property, whereas the concentration of power that occurs on the bowtie axis is a higher order property. This difficulty is avoided by using a higher order power of the Fourier coefficients, namely:

$$W(f_x, f_y, f_t) \ = \ |\hat{I}(f_x, f_y, f_t)|^2$$

Running a similar argument as above on the expected values yields:

$$\begin{aligned}
& \mathscr{E}\,\{\,|\,\hat{W}(f_x,f_y,f_t)\,|^2\,\} \\
= \; & \mathscr{E}\{\,|\,\hat{I}(f_x,f_y,f_t)\,|^4\,\} \\
= \; & \mathscr{E}\,\{\,(\,|\,\hat{I}_1(f_x,f_y,f_t))\,|^2 + |\,\hat{I}_2(f_x,f_y,f_t)\,|^2 \; + 2\,Re(\,\hat{I}_1(f_x,f_y,f_t)\,\overline{\hat{I}_2(f_x,f_y,f_t)}\,)\,)^{\,2}\} \\
= \; & \mathscr{E}\,\{\;|\,\hat{I}_1(f_x,f_y,f_t))\,|^{\,4}\,\} + \mathscr{E}\,\{\,|\,\hat{I}_2(f_x,f_y,f_t))\,|^{\,4}\;\} \\
& + \mathscr{E}\,\{\,(\,Re(\,\hat{I}_1(f_x,f_y,f_t)\,\overline{\hat{I}_2(f_x,f_y,f_t)}\;)^2\,\} + 4\,\mathscr{E}\,\{\;|\,\hat{I}_1(f_x,f_y,f_t)|^2\,|\,\hat{I}_2(f_x,f_y,f_t)\,|^2\,\} \\
& + \mathscr{E}\,\{\,Re(\,\hat{I}_1(f_x,f_y,f_t)\,\overline{\hat{I}_2(f_x,f_y,f_t)}\,)\;\;(\,|\hat{I}_1(f_x,f_y,f_t)\,|^{\,2} + |\hat{I}_2(f_x,f_y,f_t)\,|^{\,2})\,\}\,.
\end{aligned}$$

The first four terms in the summation are all positive, whereas the last term vanishes, since $Re(\,\hat{I}_1(f_x,f_y,f_t)\,\overline{\hat{I}_2(f_x,f_y,f_t)}\,)$ is equally likely to be positive as negative.

In particular, we observe the desired interaction of the two motion planes along the bowtie axis. There, the third and fourth terms in the summation are positive since both $\hat{I}_1(f_x,f_y,f_t)$ and $\hat{I}_2(f_x,f_y,f_t)$ are simultaeously non-zero and each of the terms involves squared (and hence positive) quantities only. These terms add a large extra weight to the minimization, which causes the first eigenvector e_1 of $\mathbf{F^T}\,\mathbf{W^2}\,\mathbf{F}$ to align with the bowtie axis.

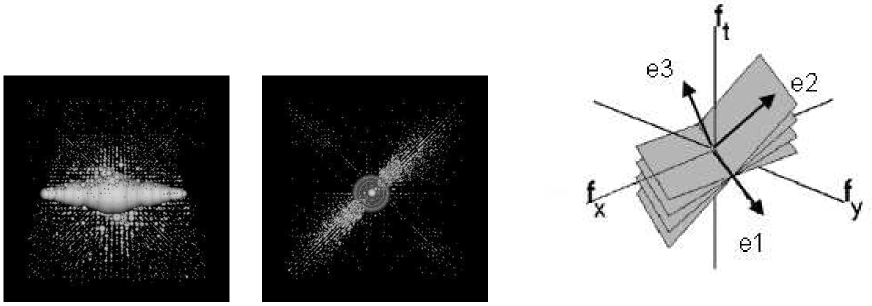

(a) (b) (c)

Figure 2: (a) Energy distribution in the Fourier domain from a viewpoint perpendicular to the axis of the bowtie (b) Energy distribution in the Fourier domain from a viewpoint parallel to the axis of the bowtie (c) Eigenbasis of bowtie signature.

3.3 Results

To evaluate our method we rendered several synthetic image sequences of scenes containing lambertian spheres (see Fig. 3). Image motion was generated by moving a camera (45° field of view) through the scene with various translation and rotation parameters. For *lateral* camera motion (axis of translation and axis of rotation both perpendicular to the optical axis) and a small field of view, the motion satisfies Eq. (1) and a bowtie should be present.

Table 1 shows the results for the PCA method and the wedge detector. Except for the real sequences in the last two rows, each row shows the average over ten runs of each method. We report both the error and the standard deviation σ of estimated bowtie parameters θ and $U = \sqrt{u_x^2 + u_y^2}$. We use both synthetic and real scenes to compare the accuracy of the PCA method to the wedge-based method of [12, 14, 18]. The results are very similar. Temporal shearing (see Sec. 2.2) is used as a preprocessing step of the PCA method to account for temporal aliasing. We also perform a power normalization [7] for each spatial frequency (f_x, f_y) to ensure uniform power distribution. Note that the PCA method is parameter-free. The wedge method needs two parameter, the slope of the wedge v_{max} and the cylinder radius r (see Sec. 2.2). For all sequences, v_{max} was set to 1.0, except for the holly sequence, where 2.0 was used. A radius r of 4 pixels was used in all cases. These values are somewhat image dependent, which emphasizes the advantage of getting rid of parameters.

The average running times for the wedge method was 30 seconds[1]. Under similar conditions, the PCA method runs in around 1.5 seconds.

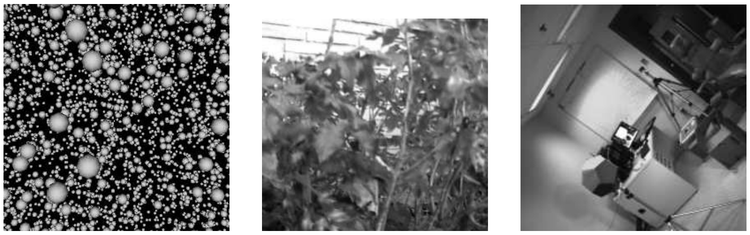

(a) (b) (c)

Figure 3: (a) Synthetic scenes were constituted of small balls at different depths (64 frames of size 128x128 pixels). (b) Holly bush taken from a horizontally translating camera (32 frames of size 64x64 pixels). (c) Lab sequence taken from a camera translating diagonally (40 frames of 128x128 pixels).

Table 1 : Comparison between the PCA method and the Wedge method

Ground truth		**PCA method**		**Wedge method**	
θ	$\lvert U\rvert$	$err(\theta)$ $(\sigma(\theta))$	$err(\lvert U\rvert)$ $(\sigma(\lvert U\rvert))$	$err(\theta)$ $(\sigma(\theta))$	$err(\lvert U\rvert)$ $(\sigma(\lvert U\rvert))$
90	0	0.24 (1.51)	0.00 (0.00)	1.16 (4.17)	0.03 (0.02)
90	0	0.58 (0.88)	0.01 (0.00)	0.59 (3.63)	0.03 (0.05)
0	0	0.33 (1.18)	0.00 (0.00)	0.14 (5.00)	0.03 (0.03)
135	0	0.08 (1.12)	0.00 (0.00)	0.35 (1.85)	0.02 (0.02)
60	0	1.48 (1.30)	0.01 (0.01)	3.92 (4.13)	0.07 (0.04)
225	0.71	0.93 (1.73)	0.02 (0.01)	0.88 (1.91)	0.02 (0.01)
90	1	2.45 (1.16)	0.05 (0.02)	3.13 (5.09)	0.06 (0.04)
135	1.41	1.11 (2.83)	0.04 (0.01)	0.04 (2.53)	0.04 (0.03)
Real Sequence		θ	$\lvert U\rvert$	θ	$\lvert U\rvert$
Holly		92.78	0.01	88.95	0.16
Lab		28.58	0.02	26.01	0.05

3.4 Comparing the eigenvalues

Optical snow occurs for scenes with a wide range of depths. However, a wide range of depths does not necessarily produce a bowtie with a wide range of speeds. There are

[1] The machine was an AMD Athlon 1600MHz.

several reasons why such a failure can occur. First, if the objects are all distant (a mountain seen by a walking observer) then the translation component of the image velocities vanishes, since the translation component depends on inverse depth [16]. Second, even if the objects are nearby, if the observer is moving directly forward and the field of view is small, then the translation component of image motion is small. The reason is that the translation component varies directly with the angular distance from the direction of heading. Third, if the camera motion is pure rotation (no translation), then the image motion field is independent of depth. In all three of these cases, the "bowtie" in fact consists only of a single motion plane.

When only a single motion plane is present, it is the best fitting plane to the "bowtie". This plane π is spanned by the eigenvectors e_1 and e_2. The eigenvalues would depend on the spatial orientation structure in the (translating) image, rather than on a bowtie axis. Moreover, the eigenvalue for the e_3 direction would be near zero.

With this issue in mind, we would like to estimate how well the analyzed motion corresponds to optical snow, i.e. if a bowtie is present. We define a bowtie "fitness measure" by comparing the two largest eigenvalues. There are two extreme cases to consider. The first is that optical snow is present and there is a wide range of image speeds. In this case, the first two eigenvalues should correspond to the bowtie axis and a perpendicular vector in the best fitting plane π, respectively. The second extreme case is that a single motion plane is present. Assuming that spatial content of the image is roughly evenly distributed over all orientations, the two corresponding eigenvalues should be equal.

Figure 4 shows a plot of the ratio of the first two eigenvalues λ_2 and λ_1 as a function of the range of depths. As predicted, the ratio falls off as the range of depths decreases.

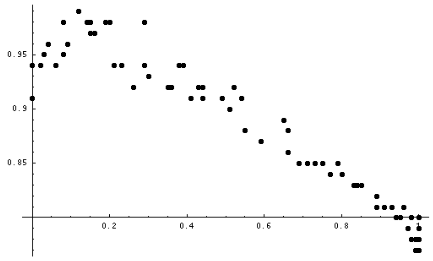

Figure 4: $\frac{\lambda_2}{\lambda_1}$ as a function of depth range (from a single depth to a large interval). As expected, it starts at 1 (pure plane) and decreases as a bowtie signature takes shape.

4 Conclusion

We presented a new simple PCA-based method to analyze optical snow. The performance of the method in estimating the bowtie axis is similar to that of the previous method, which is based on a wedge filter. As the results demonstrated, the PCA method has many advantages such as its simplicity, its efficiency and its absence of parameter.

References

[1] J.L. Barron, D.J. Fleet, and S.S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, February 1994.

[2] J.R. Bergen, P.J. Burt, R. Hingorani, and S. Peleg. A three-frame algorithm for estimating two-component image motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(9):886–896, September 1992.

[3] M. J. Black and D. J. Fleet. Probabilistic detection and tracking of motion discontinuities. In *International Conf. on Computer Vision*, pages 551–558, Corfu, Greece, Sept. 1999.

[4] M.J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow-fields. *Computer Vision and Image Understanding*, 63(1):75–104, January 1996.

[5] T. Darrell and A.P. Pentland. Cooperative robust estimation using layers of support. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):474–487, May 1995.

[6] N.M. Grzywacz and A.L. Yuille. A model for the estimate of local image velocity by cells in the visual cortex. *Proceedings of the Royal Society of London. B*, 239:129–161, 1990.

[7] D.J. Heeger. Optical flow from spatiotemporal filters. In *First International Conf. on Computer Vision*, pages 181–190, 1987.

[8] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[9] C.L. Huang and Y.T. Chen. Motion estimation method using a 3d steerable filter. *Image and Vision Computing*, 13:21–32, 1995.

[10] M. Irani, B. Rousso, and S. Peleg. Computing occluding and transparent motions. *International Journal of Computer Vision*, 12(1):5–16, February 1994.

[11] A.D. Jepson and M.J. Black. Mixture models for optical flow computation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 760–761, 1993.

[12] M. S. Langer and R. Mann. Dimensional analysis of image motion. In *IEEE International Conference on Computer Vision*, pages 155–162, 2001.

[13] M. S. Langer and R. Mann. Tracking through optical snow. In *Lecture Notes in Computer Science 2525: Biologically Motivated Computer Vision, 2nd International Workshop*, pages 181–188. Springer, November 2002.

[14] M.S. Langer and R. Mann. Optical snow. *International Journal of Computer Vision*, 55(1):55–71, 2003.

[15] M. Lappe and J. P. Rauschecker. A neural network for the processing of optical flow from egomotion in man and higher mammals. *Neural Computation*, 5:374–391, 1993.

[16] H.C. Longuet-Higgins and K. Prazdny. The interpretation of a moving retinal image. *Proceedings of the Royal Society of London B*, B-208:385–397, 1980.

[17] B.D. Lucas and T. Kanade. Optical navigation by the method of differences. In *International Joint Conf. on Artificial Intelligence*, pages 981–984, 1985.

[18] R. Mann and M. S. Langer. Estimating camera motion through a 3d cluttered scene. In *Canadian Conference on Computer and Robot Vision*, London, Canada, May (to appear) 2004.

[19] K.M. Mutch and W.B. Thompson. Analysis of accretion and deletion at boundaries in dynamic scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(2):133–138, March 1985.

[20] H.H. Nagel. On the estimation of optical flow: Relations between different approaches and some new results. *Artificial Intelligence*, 33(3):299–324, November 1987.

[21] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11):559–572, 1901.

[22] E. P. Simoncelli. *Distributed Analysis and Representation of Visual Motion*. PhD thesis, Massachusetts Institute of Technology, 1993. Department of Electrical Engineering and Computer Science.

[23] A. Watson and A. Ahumada. Model of human visual-motion sensing. *Journal of the Optical Society of America*, 2(2):322–342, 1985.

[24] Y. Weiss. Smoothness in layers: Motion segmentation using nonparametric mixture estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 520–526, 1997.