Image Analysis For Avalanches

Jim McElwaine

Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, Silver Street, Cambridge, CB3 9EW.

jnm11@amtp.cam.ac.uk

ABSTRACT: *An algorithm for calculating the motion of avalanches from films based on changepoint determination is described. The method is designed for the detection of moving objects against a static background and works for partially transparent objects and objects with spatially and temporarily varying colour components. The method can be applied to mono-spectral or multi-spectral signals and works well even with low quality video pictures and uncontrolled illumination. The method is particularly suitable for tracking granular flows and other scientific experiments.*

KEYWORDS: *Image analysis, motion detection.*

Table of Contents

1 Introduction

Analysing films of avalanches can provide information about their position, size and velocity. The data is invaluable for testing and developing theories and provides much more information than almost any other sort of sensor. The data can be directly compared with integral (lumped mass) models or continuum models.

However, analysing avalanche film is difficult and this has rarely been done quantatively. Current methods mostly involve an operator going through a film picture by picture and marking the outline of the avalanche by hand. Some automation has been achieved [LAT 97] when the boundary of the avalanche was marked in one picture and then tracked by computer. Any method that requires substantial input from a user introduces a subjective element, different users will pick different edges, and the same user will pick a slightly different edge if they repeat the analysis. This subjectivity means accurate error estimation is impossible. If many films are to be analysed it is also very time consuming if substantial user intervention is required.

This paper describes a method that is fully automatic which can be applied to a wide variety of flows. Since films of avalanches, or indeed any flow experiment, are expensive compared to the cost of computing power, we tried to develop an off-line algorithm that made maximal use of the data without regard to the computational cost.

The algorithm described in this paper was designed to analyse the video footage from a series of experiments in Japan to investigate the structure of avalanches [NIS 98, MCE 01]. The experiments took place at the Miyanomori Ski Jump in Japan (the normal hill from the 1972 Sapporo Olympics). In these experiments up to 650,000 ping-pong balls were released at the top of the landing slope of the Miyanomori ski jump and allowed to flow down the slope in an avalanche. In 1999 around eighty clips of film were taken from varying positions and this algorithm was developed to analyse these. The results of this analysis are presented in [MCE 02].

2 Digital Image Processing

Digital image processing is an enormous field [PRA 91]. One of the most basic aims is the identification of objects in image sequences. This requires a model of the object and the background. The simplest case is when they have known colours and an image pixel can be classified by assigning them to the *closest* colour, using some suitable metric. More usually the exact colours are not known in advance and they are estimated using statistics calculated from the pictures. This technique extends naturally to the case of image sequences where now the segmentation takes place in a three dimensional space [DEN 99, LIM 90, JÄH 97]. Other difficulties can be caused by uneven illumination between the top and bottom of the slope, and by shadows. These can cause the same object to have a different colour when it is on different parts of the screen.

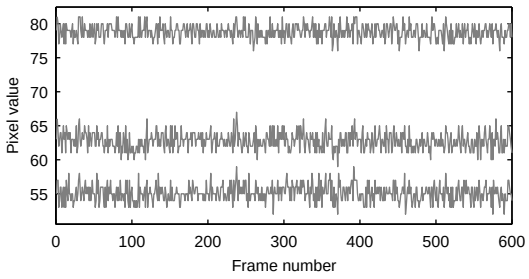

Figure 1: *RGB (Red, Green Blue) values for a background pixel.*

Many objects and backgrounds are not of uniform colour and more complicated techniques are necessary. If the object and background have slowly varying colours relative to the difference between them then edge detection techniques can be used to calculate the boundaries and thus the regions. More sophisticated techniques are not applied directly to pixel value but to derived quantities such as temporal or spatial gradients, or structure tensors that are constructed so as to include texture information.

If information about the scene is incorporated into the algorithm more powerful techniques can be used. For example, if it is known that the scene can be divided into regions separated by simple curves then these curves can be tracked from picture to picture. This approach sometimes called snakes or active contours [KAS 87, MA 97]has been applied to snow avalanches [LAT 97]. This approach works well when there is a simple difference in some spatial property of the image between the objects. For example they have different colours or textures.

If the experimenter has full control over the scene in his picture they can film against a uniform background that has high contrast compared to the objects he wishes to identify, and the lighting can be arranged so as to ensure that there are no shadows. Under such controlled conditions the previously mentioned approaches work very well. In other cases when the film is taken under uncontrolled conditions, such as outside by amateur cameramen, this is not possible. With snow avalanches there is the obvious difficulty that an avalanche is snow moving on a background of snow. The difference in intensity and colour between the avalanche and the background is very small. However, there are differences in texture between the background and the avalanche, but these are hard to quantify.

3 Detecting First Arrival Times

Algorithms based on an accurate underlying model are more powerful than those based on less information. The better the model describes and restricts the data the better the algorithm will work and the more accurate the results will be. The previously mentioned techniques add very little information, though the active contour

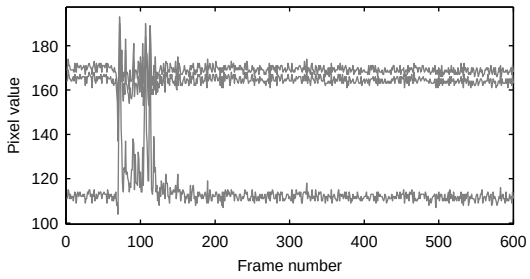

Figure 2: *RGB values when the avalanche arrives around frame number 80.*

approach does introduces smooth motion of the avalanche front from frame to frame. In particular all of these approaches are local in time, that is they only consider a few pictures at a time. An obvious feature of avalanches is that they move downhill. This means that any point on the slope will have a roughly constant colour until the avalanche arrives. It will have a varying colour while the avalanche flows over that point of the slope. Once the avalanche has flowed past the colour will again be approximately constant, though perhaps a different constant. Figure 1 shows a pixel from a position on the slope where the avalanche never arrives. Figure 2 shows a pixel from a position on the slope where the avalanche arrives around frame 70 and leaves around frame 120. It is very obvious from these signals if and when the avalanche arrives, even if the flowing material is the same colour as the background, because the signal within the avalanche has much greater variation. If the first arrival time can be calculated for all the pixels then the front of the avalanche at any time is represented by a contour at that time.

Separating a signal into regions with different statistical properties is known as *changepoint detection* and is a well studied problem in statistics [BAS 88, BAS 93, ÓR 96]. We are concerned with *off-line* changepoint detection, that is using the whole signal to determine where the changes occur.

We consider two approaches to choosing the changepoints MLE (Maximum likelihood estimation) and the Bayesian approach MAP (Maximum A priori Probability), and show that the results are largely equivalent.

4 Maximum Likelihood Estimation

The likelihood of a dataset x is it's probability of occurring $p(x)$. In Bayesian analysis this is synonymous with evidence. We consider a vector valued signal $\boldsymbol{x}$ at k times $\{\boldsymbol{x}_j \in \mathbb{R}^d : j = 1, \ldots, k\}$. We suppose that the signals are drawn from normal distributions with unknown mean values $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariances $C_i \in \mathbb{R}^{d \times d}$. First we consider a signal where all elements are identically distributed. Later we consider signals where the parameters $\boldsymbol{\mu}$ and C are different in regions separated by changepoints.

The probability that $\{\boldsymbol{x}_j : j = 1 \ldots k\}$ are drawn from $N(\boldsymbol{\mu}, C)$ is

$$P(\{\boldsymbol{x}_j : j = 1 \ldots k\}|\boldsymbol{\mu}, C) = \frac{\exp\{-\frac{1}{2}\sum_i^k (\boldsymbol{x}_i - \boldsymbol{\mu})C^{-1}(\boldsymbol{x}_i - \boldsymbol{\mu})\}}{(2\pi)^{\frac{1}{2}kd}|C|^{\frac{1}{2}k}}, \tag{1}$$

where $|C|$ is the determinant of C. If we define $\hat{\boldsymbol{\mu}} = \frac{1}{k}\sum_i \boldsymbol{x}_i$ the sample mean and $\hat{C} = \frac{1}{k}\sum_i (\boldsymbol{x}_i - \hat{\boldsymbol{\mu}})(\boldsymbol{x}_i - \hat{\boldsymbol{\mu}})^T$ the sample covariance. This can be written

$$P(x|C, \boldsymbol{\mu}) = \frac{\exp\{-\frac{1}{2}k(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})^T C^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) - \frac{1}{2}k\,\mathrm{Tr}(\hat{C}C^{-1})\}}{(2\pi)^{\frac{1}{2}kd}|C|^{\frac{1}{2}k}}. \tag{2}$$

In the Generalised Maximum Likelihood Ration (GMLR) approach the unknown parameters $\boldsymbol{\mu}$ and C are replaced by their maximum likelihood estimations, which are

$$\frac{\partial \log P}{\partial \boldsymbol{\mu}} = kC^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) = 0 \quad \Rightarrow \quad \boldsymbol{\mu} = \hat{\boldsymbol{\mu}}, \tag{3}$$

$$\frac{\partial \log P}{\partial C^{-1}} = \tfrac{1}{2}kC - \tfrac{1}{2}k(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu})^T - \tfrac{1}{2}k\hat{C} = 0 \quad \Rightarrow \quad C = \hat{C}. \tag{4}$$

Substituting this back into eq. 2 gives the log likelihood as

$$\log P = -\tfrac{1}{2}kd\log(2\pi e) - \tfrac{1}{2}k\log|\hat{C}| \tag{5}$$

In the Bayesian framework we regard $\boldsymbol{\mu}$ and C as random variables with specified prior distributions [O'H 99, BER 94]. The standard prior distributions are to take $\boldsymbol{\mu}$ normal with mean $\boldsymbol{\nu}$ and covariance C/β, and C^{-1} to have a Wishart distribution with exponent α and trace $T\alpha$. The log of the *evidence* $P(x)$ is then

$$\begin{aligned} \log P(x) &= \int (x|C, \boldsymbol{\mu})p(\boldsymbol{\mu}|C)P(C)\, d\boldsymbol{\mu}\, dC \\ &\approx \text{const} - \tfrac{1}{2}kd\log(2\pi e) - \tfrac{1}{2}(k+\alpha)\log|\hat{C}'|, \end{aligned} \tag{6}$$

The updated covariance

$$\hat{C}' = \frac{\alpha}{k+\alpha}T + \frac{k}{k+\alpha}\hat{C} + + \frac{k}{k+\alpha}\frac{\beta}{k+\beta}(\hat{\boldsymbol{\mu}} - \boldsymbol{\nu})(\hat{\boldsymbol{\mu}} - \boldsymbol{\nu})^T. \tag{7}$$

This is essentially the same likelihood as the GMLR calculation eq. 6, but with the prior distributions acting like additional data. The prior for $\boldsymbol{\mu}$ is equivalent to β data samples having mean $\boldsymbol{\nu}$, and the prior for C acts like α data samples with covariance matrix T. If we wish to add additional data to the model, for example information that the covariance of background pixels is small, this can easily be included by choosing the prior variables. By taking limits as $\alpha \to \infty$ a known covariance can be included or $\beta \to \infty$ a known mean. For the ping-pong ball avalanches the algorithm performed well without using using any prior information that is $\alpha = 0$ and $\beta = 0$. For a film of a natural snow avalanche from Ryggfonn using prior information was useful and is discussed later.

5 Changepoint Likelihood

Consider a sequence of n data points such that the first k are drawn from one normal distribution and the next $n-k$ from a different one. Let $p(k)$ denote the likelihood of this changepoint then using eq. 6

$$\begin{aligned}\log p(k) &= \log p(\{\boldsymbol{x}_i | i = 1 \dots k\}) + \log p(\{\boldsymbol{x}_i | i = k+1 \dots n\}) \\ &= -\tfrac{1}{2} n d \log(2\pi e) - \tfrac{1}{2} k \log |\hat{C}_{1\dots k}| - \tfrac{1}{2}(n-k) \log |\hat{C}_{k+1\dots n}|.\end{aligned}$$

Where $\hat{C}_{j\dots k}$ denotes the covariance of $\{\boldsymbol{x}_i | j \le i \le k\}$. To find the changepoint this function must be calculated for all k and the maximum value chosen.

This can be calculated efficiently as follows with a number of operations that scales as nd^2. A simple approach, followed in this paper for the ping-pong ball avalanches, is to take the data to be the RGB (Red, Green and Blue) values from each pixel so that $d = 3$. For the Ryggfonn snow avalanches there is very little colour in the pictures and nearly identical results were obtained using RGB values or the sum of these so that $d = 1$. There are many other possible choices for the data $\{\boldsymbol{x}_j\}$, such as transforming the pixel values or combining data from more than one pixel. A transform to compensate for changes in exposure is discussed in section 7. A desirable property of this algorithm is that it is coordinate invariant. That is, identical results would be obtained if the data were transformed by a rotation, for example YUV space, which is the native data space for many video recording formats. Define the cumulative sums

$$\boldsymbol{m}_j = \sum_{i=1}^{j} \boldsymbol{x}_i \qquad s_j = \sum_{i=1}^{j} \boldsymbol{x}_i \boldsymbol{x}_j^T \qquad [8]$$

Then

$$\begin{aligned}C_{1\dots k} &= \frac{s_k}{k} - \frac{\boldsymbol{m}_k \boldsymbol{m}_k^T}{k^2} \\ C_{k+1\dots n} &= \frac{s_n - s_k}{n-k} - \frac{(\boldsymbol{m}_n - \boldsymbol{m}_k)(\boldsymbol{m}_n - \boldsymbol{m}_k)^T}{(n-k)^2}\end{aligned}$$

$p(k)$ can thus be calculated by first going through the data calculating the total sum values $\boldsymbol{m}_n$ and s_n. Then making another pass where the cumulative sums are updated, the likelihood calculated, and then the maximum found. We use a slightly different function to maximise so that it is normalised, scale invariant and always positive.

$$M(k) = \log |\hat{C}_{1\dots n}| - \frac{k}{n} \log |\hat{C}_{1\dots k}| - \frac{n-k}{n} \log |\hat{C}_{k+1\dots n}|. \qquad (9)$$

$M(k)$ measures the increase in the likelihood by introducing a changepoint at k as opposed to modelling the whole signal as being drawn from the same distribution.

There are two small technical issues that need to be addressed. Firstly the pixel data is quantised and this reduces the covariance. If we assume that the true values

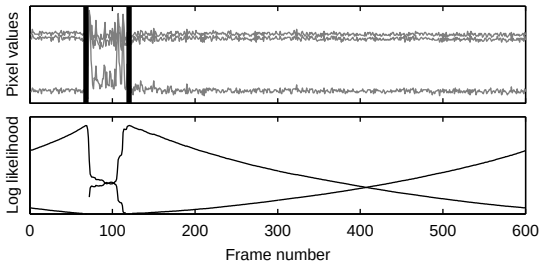

Figure 3: *Multiple Changepoints, pixel values and log-likelihood.*

are independently and uniformly distributed over the quantisation interval, then the reduction in covariance is $\frac{1}{6}I$, where I is the identity matrix. Compression algorithms can introduce larger quantisation errors so we include this as one of the model parameters and add λI to each covariance before calculating the determinant. If λ is taken to be large then

$$\log|C + \lambda I| = \mathrm{Tr}\log(C + \lambda I) = d\log\lambda + \mathrm{Tr}\, C + O(1/\lambda).$$

Therefore

$$M(k) \approx \mathrm{Tr}\,\hat{C}_{1\ldots n} - \frac{k}{n}\,\mathrm{Tr}\,\hat{C}_{1\ldots k} - \frac{n-k}{n}\,\mathrm{Tr}\,\hat{C}_{k+1\ldots n}.$$

Thus $M(k)$ does not depend on λ and now is a function of the mean squared deviations $\mathrm{Tr}\, C$. In some situations this may give better results, But in general small values of λ are better since the algorithm is more sensitive to changes in the different channels. For example, if the green value is constant and then changes to a new constant value at the changepoint this would give zero determinants, thus infinitely high likelihood, for the correct changepoint regardless of noise in the other channels. Whereas the high λ limit will only detect changes which are large in mean square summed over all channels.

Secondly introducing a small value for λ also deals with another numerical issue that the determinants will be zero, and thus the logs infinite, unless there are more than $(d+3)/2$ samples. If uninformative priors had been used this would have been reflected in factors of $k-3$ and $n-k-3$ (when $d=3$) in front of the logs in eq. 9 making the likelihood zero or *minus infinity* for fewer than four samples.

6 Multiple Changepoints

So far we have only discussed the presence of one changepoint, but in most cases there are two that must be detected, that is when the avalanche arrives and when the avalanches leaves. In the case of the ping-pong balls the background will be the same after the balls have passed, so that we search for two changepoints that segment

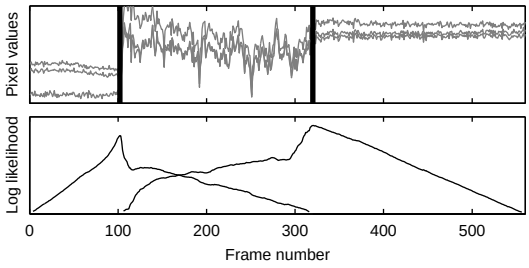

Figure 4: *Multiple Changepoints, pixel values and log-likelihood with changed background. The front arrives around frame 100 and comes to a halt around frame 300.*

the data into two statistically homogeneous regions. For snow avalanches, or regions where the balls become stationary, the background mean value will have changed after the avalanche has passed or stopped moving so that the two changepoints segment the data into three statistically homogeneous regions. In either case the optimisation now has to be carried out over two variables, the two changepoints, rather than the one. A direct search would be expensive (order n^2 operations) and is not necessary. Instead we alternate optimising the changepoints individually holding the other fixed until convergence occurs. Typically this is extremely rapid and involves only one or two searches for each changepoint. The top plot of fig. 3 shows a data sample and the bottom plot shows the log likelihoods for the changepoints with the other held fixed. The maximum of these functions clearly occurs at the arrival and exit of the avalanche. Fig. 4 is similar to fig. 3 but there balls come to rest so that the background has changed between the start and the end of the film. The extra variation in the signal between the two changepoints is clearly visible. This could be used as the start of a method of calculating internal velocities.

The searching for the maximum likelihood changepoints is a global method that is very robust and nearly always locates the correct changepoints to within a few frames. However, there are other properties of the signal that have not so far been included here that can reduce the accuracy of the changepoint estimations. These can all be regarded as a failure in the model, which is that the pixel values are unchanging before the avalanche arrives. There are at least three causes of this. Compression of video signals and band-limiting can introduce ringing and smearing between adjacent pixels, so that a pixel value changes in advance of the avalanche arrival. This can be minimised by using high quality video equipment, and working with video that has been compressed without loss. Most video compression algorithms are designed to exploit the perceptual properties of the human vision system, such as reduced chroma resolution. Even though there can be negligible perceptual loss this information loss will degrade the performance of any image processing algorithm and should be avoided by working with lossless, or nearly lossless, compression systems. Other causes are shadows which appear as a darker edge moving in front of the flow, and overall changed in the intensity of all the pixels as the camera adjusts its exposure in response to the changing scene. These last two effects are discussed in the following two sections.

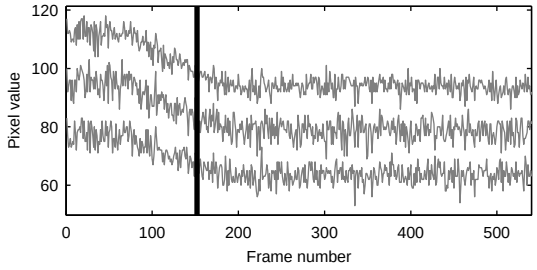

Figure 5: *Changing pixel values due to the camera's auto-exposure producing false changepoints.*

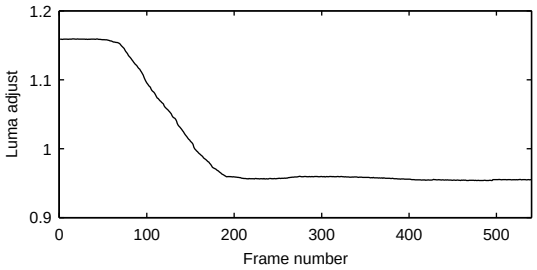

Figure 6: *The scale factor necessary to adjust for changes in exposure.*

7 Luminance Adjustment

If the illumination of the scene changes the pixel values will be changed and this can cause the determination of incorrect changepoint. This effect is rarely important over film clips of a few seconds, but changes in the cameras exposure settings are very important. The details of this will vary from camera to camera, but broadly speaking if auto-exposure is operating a camera will adjust the exposure, aperture and shutter speed, so as to maintain a constant brightness in the image. If a large white object moves into the field of view, in our case a ping-pong ball avalanche, the camera will reduce the exposure and the background pixels will darken. This problem would of course be removed by turning off auto-exposure, but since there are many films where this was not done it is necessary to account for it. This can be clearly seen in fig. 5 where a changepoint is incorrectly detected.

If we consider initially only background pixels which we index by x, where the index is over position and colour components, and by t which indexes frame number, we can consider the complete image sequence as a function $f(x,t)$ where t is the frame number and x the pixel coordinates. This can be written

$$f(x,t) = f_b(x)f_t(t) + Z(x,t), \tag{10}$$

where $f_b(x)$ is the underlying pixel value at x, f_t is the illumination and $Z(x,t)$ is a

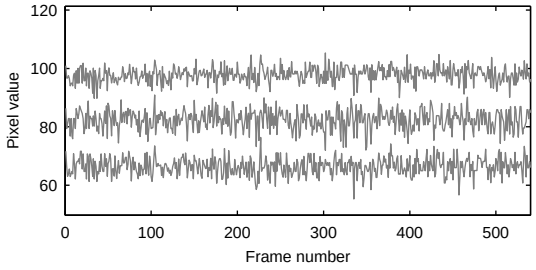

Figure 7: *The same pixel values after exposure compensation.*

noise term.

This model assumes that video system is linear in intensity. If the nonlinear response of the video system, for example the gamma curve, was known the data could first be linearised using this information and the same analysis made. Another possibility would be separate light intensities f_t for Red, Green and Blue, if the colour balance changes. For the films analysed in this paper neither of these refinements were necessary.

A maximum likelihood approach is equivalent to minimising the squared error

$$\sum_{x,t}[f(x,t) - f_b(x)f_t(t)]^2.$$

The solutions of this are the left and right eigenvectors corresponding to the largest singular value of $f(x,t)$ considered as a matrix.

$f(x,t)$ is a large matrix typically of order 1,000,000 by 500. Fortunately the largest singular value is much larger than the next, provided the background model is accurate. Thus direct iteration gives convergence to machine precision in around 10 iterations

$$f_b(x) \to \sum_t f(x,t)f_t(t), \quad f_t(x) \to \sum_x f(x,t)f_b(x)$$

The iteration is started by taking $f_b(x)$ to be the mean background picture and $f_t(t)$ to be constant.

Parts of the scene are not constant and would distort the luminance correction. The iteration scheme is modified so that only pixels with a root mean squared deviation of less than a threshold value 5 are included in the iteration. The value 5 was found to work well for all the cases considered, because it is just above the standard deviation of most background pixels. This threshold could instead be chosen dynamically if it varied a lot from film to film. This non-linearity means that convergence is no longer guaranteed, but in all cases considered convergence was achieved in no more than 20 iterations. This approach has the useful feature of providing a very robust segmentation of the image sequence into the background region, and the region which

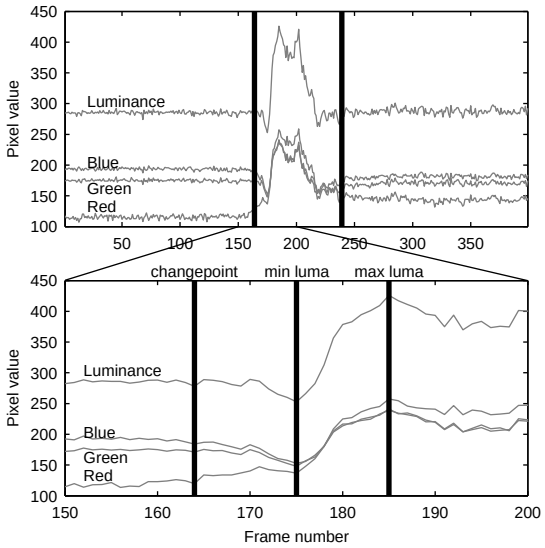

Figure 8: *The lower figure is a close up of the upper figure and shows shadow correction moving the first changepoint from frame 164 to 175. The luminance is the Euclidean length of the colour triple.*

the avalanche enters. Figure 6 shows the calculated luminance function $f_t(t)$. This function can then be applied to all pixels to remove the changing exposure effect. This is shown in fig. 7 which is the same data as fig. 5, but after the correction has been applied.

If luminance correction is not required, thresholding into regions where the flow enters and regions never entered can be achieved more quickly by calculating the covariance of each pixel over the whole image sequence and then segmenting this.

8 Shadow Adjustment

The flow can influence pixels before it reaches a point by casting shadows illustrated in fig. 8. The overall luminance drops significantly in front of the avalanche head, even though the red component rises a little, and this causes the changepoint to be too early. It is typically the case that the changepoint will be chosen too early as the values within the flow have much higher variance than the background. Different strategies for finding the shadow are possible. A perfect shadow with no noise would result in a monotonic decrease in all three colour components. So one could search forward as far as all three are decreasing. This works in some cases but would fail for fig. 8 as the red channel is increasing. Another choice would be to search forwards as long as the total luminance, the square root of the sum of all three colour components, is decreasing. Again in some cases this works but fails in fig. 8 as noise in the signal

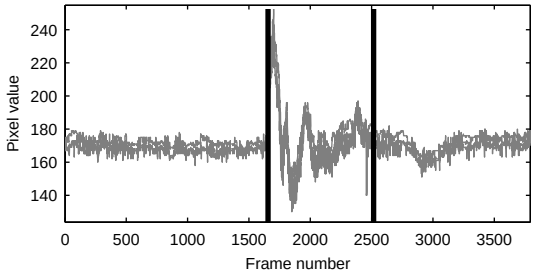

Figure 9: *Time slice from an avalanche in Ryggfonn. The pixel is in the centre of the avalanche and is not in shadow. The changepoint marked around 1600 correctly detects the arrival of the avalanche around frame 1600. The tail of the avalanche is less clear.*

means that the luminance is not monotonic decreasing. A more robust approach is to look for a global minimum of the luminance on some restricted region. Since the ping-pong balls are brighter than nearly all of the background and the signals typically start with a peak we first search forwards a fixed distance looking for the largest luminance value. Then a search is carried out for the minimum luminance between the original changepoint and the maximum luminance. The results of this procedure are shown in the bottom plot of fig. 8. Typical window widths for shadow searching will depend on the nature of the flow; but for the data considered here 20 worked well.

These methods are rather ad hoc. A more general method in keeping with the approach of this paper would be to model statistically the shadow region, and find changepoints on either side. A possible model might be

$$\boldsymbol{x}_i = \alpha_i \boldsymbol{\mu} + X_i, \tag{11}$$

which is the same as the model used for luminance correction. X_i is a noise term with zero mean and the same covariance as the background, and $\boldsymbol{\mu}$ is the background mean value. $1 \geq \alpha_i \geq 0$ are unknown nuisance parameters to be integrated out using a Bayesian approach, or eliminated using Maximum Likelihood estimation. Further conditions could be imposed on the α_i such as linear decrease $\alpha_i = 1 - (i-1)\beta$, or power decrease $\alpha_i = \beta^{1-i}$. Another approach is discussed in the next section.

9 Snow Avalanches

The algorithm has recently been applied to snow avalanches from Ryggfonn and Col du Lautaret. This work is not complete but preliminary results are discussed. The algorithm works without modification, but there are difficulties. Fig. 9 shows a time slice where the edge of an avalanche is clearly located and the algorithm works correctly. Fig. 10 shows a more difficult case from the same avalanche along the left

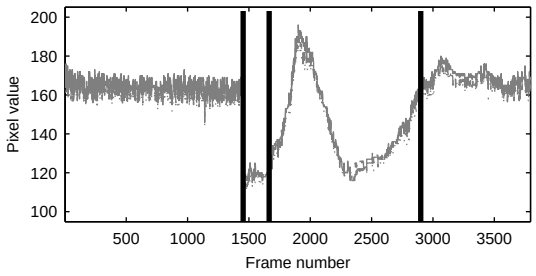

Figure 10: *Time slice from an avalanche in Ryggfonn. The pixel is on the left edge of the avalanche and is in heavy shadow. The changepoint marked around 1450 does not correctly detect the arrival of the avalanche but instead marks the arrival of the shadow.*

hand edge where there is strong shadow. The first changepoint shown in the figure is not the arrival of the avalanche but the arrival of the shadow. The second changepoint at frame 1667 is in fact the arrival of the avalanche. These changepoints are calculated by modelling the signal as four regions. The first region is the background, the second region a shadow the third region the avalanche and the fourth region the background. For the background regions the covariance from the first background region is used as prior information for the second and fourth regions enforcing low covariance.

10 Interlace Correction

The result of the changepoint algorithm with possible shadow corrections is a first arrival time for the flow at each pixel. Many video systems including PAL and NTSC are *interlaced*, that is a frame is divided into two fields of odd and even scan lines. Care needs to be taken because depending on the video system and digitisation procedure odd or even fields can be first. The interlacing can then easily be taken into account by shifting the first arrival times on the odd or the even scan lines by one field.

11 Filtering First Arrival Times

The first arrival times for each pixel are quantised by the frame rate 1/50 s for PAL or 1/60 s for NTSC. That is the true first arrival could have been anytime between one field earlier and one field later. The next step in the algorithm is to estimate this exact arrival time using local smoothing and global constraints on the flow.

Let $t(x, y)$ be the first arrival time at point (x, y),or infinity if the point is never reached. Then we wish to find a smooth field f so that $|f(x, y) - t(x, y)| < s$, where s is

the time between fields (1/50 s or 1/60 s). One physically based criteria of smoothness would be to try to minimise the implied accelerations. The velocity of the front is

$$\boldsymbol{v} = \frac{\boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|^2}, \tag{12}$$

so the acceleration is

$$\boldsymbol{a} = (\boldsymbol{v} \cdot \boldsymbol{\nabla})\boldsymbol{v} = \frac{\boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|^2} \cdot \boldsymbol{\nabla} \frac{\boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|^2}. \tag{13}$$

Since these are non-linear in f it is quite difficult to minimise these directly. Instead we use an approach that will give local smoothing and impose global constraints.

The global constraint that we impose is that there can be only one interior minima of f and no interior maxima. This is because the front of the avalanche is always advancing and because we set the arrival time for points where the avalanche doesn't reach to infinity. The one minima corresponds to the source of the avalanche. It is a great advantage of using first arrival times that these global constraints on the motion can be easily imposed. Solutions of Laplaces equation with Neumann boundary conditions have no interior turning points and provide a smooth interpolation of the boundary values. We therefore use the following regularisation which gives rise to a Laplace equation with Neumann boundary conditions:

$$\text{minimise} \iint |\boldsymbol{\nabla} f|^2 \, dx \, dy \quad \text{subject to } I(x,y)|f(x,y) - t(x,y)| < 1, \tag{14}$$

where $I(x,y) = 1$ if the data at point (x,y) is trusted and $I(x,y) = 0$ otherwise. All the boundary points are set to a number larger than the largest arrival times and fixed as trusted. The purpose of $I(x,y)$ is to prevent incorrect points from effecting the results. A point can be marked as incorrect if no changepoint was detected or if it is a local minima or maxima. One local minima is fixed as trusted corresponding to the starting point of the avalanche. Solving eq. 14 results in Laplaces equation $\boldsymbol{\nabla}^2 f = 0$ with boundary conditions specified by the active constraints. For the boundary conditions one point is fixed as a minimum of $f(x,y)$ where the avalanche enters. The outer boundary is all set to a number much larger than the largest arrival times.

The equations can be solved efficiently using successive over relaxation [GOL 89]. One iterative scheme is

$$\begin{aligned}
f(x,y) &\rightarrow \omega[f(x+1,y) + f(x-1,y) + f(x,y+1) + f(x,y-1)] \\
f(x,y) &\rightarrow \begin{cases} f(x,y) & I(x,y) = 0 \text{ or } |f(x,y) - t(x,y)| < 1 \\ t(x,y) - 1 & I(x,y) = 1 \text{ and } f(x,y) < t(x,y) - 1 \\ t(x,y) + 1 & I(x,y) = 1 \text{ and } f(x,y) > t(x,y) + 1 \end{cases} \\
I(x,y) &\rightarrow \begin{cases} I(x,y) & |f(x,y) - t(x,y)| < 1 \\ 0 & f(x,y) = t(x,y) + 1 \text{ and } f(x,y) \text{ is local maximum} \\ 0 & f(x,y) = t(x,y) - 1 \text{ and } f(x,y) \text{ is local minimum,} \end{cases}
\end{aligned}$$

where ω is the relaxation parameter. We use the optimal relaxation parameter for an $n \times m$ rectangular region with boundary conditions on the edge, which is

$$w = \frac{4}{2 + \sqrt{4 - [\cos(\pi/m) + \cos(\pi/n)]^2}}. \tag{15}$$

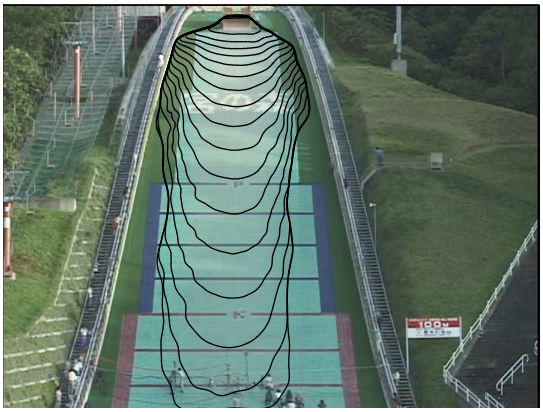

Figure 11: *Contours showing the position of the front at adjacent fields.*

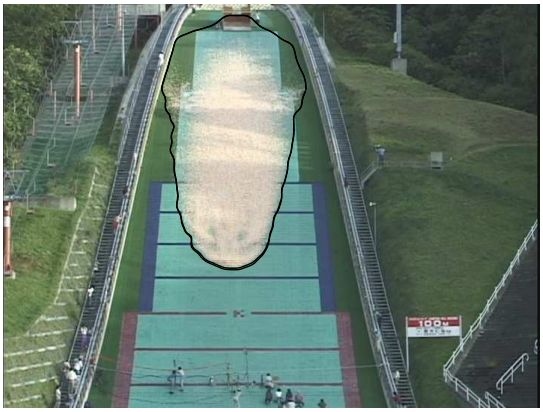

Figure 12: *Contours showing the position of the front at adjacent fields, superimposed on the avalanche from the same fields.*

The three steps are followed in turn for each position and the updated values used for the next point. The scheme is iterated until the results are changing less than a threshold. The scheme can be proved to converge.

The regularised first arrival times can then be used to plot contours of the front position as in fig 11. Because there are no local minima or maxima except at the start position there is exactly one contour for each time. Figure 12 shows two contours from adjacent fields superimposed on the actual flow. Close examination of the data shows that the contours are accurate to within a pixel near the front.

12 Filming Avalanches

This section contains a few suggestions for filming avalanches so as to maximise the accuracy of subsequent analysis.

Turn the camera on its side. Video cameras typically have horizontal resolution of around 720 pixels and a vertical resolution of 525 or 625. However, because of interlacing the full vertical resolution is only available at half the sampling frequency. Therefore if the main flow direction is horizontal in the picture frame the accuracy of velocity measurements can be doubled.

Avoid shadows. When possible shoot with the sun, or light source, behind the camera by choosing the camera location and time of day carefully. Shadows and other effects such as glare, can be dealt with but increase the complexity of algorithms and can reduce accuracy.

Set camera to manual. If the camera alters the focus, white balance, or exposure during the image sequence the analysis is more complicated. It is also important to keep the camera as still as possible.

Choose a long exposure. A long exposure time corresponds to high pass filtering and reduces under-sampling aliasing errors. Ideally the exposure time would be the same as the field rate.

13 Conclusions

This paper has presented a new approach to analysing image sequences from avalanches. The method has wide applicability and should be useful in many fields. From the smoothed first arrival times it is easy to calculate the velocity and acceleration of the avalanche edge at any time using eq. 12 and eq. 13. This provides valuable data for testing and developing avalanche theories. The main limitations of this approach are that the camera must be fixed and unzoomed throughout the sequence.

14 Acknowledgments

This work was carried out while the author was a European Union Research Fellow at the Institute of Low Temperature Science, Hokkaido University, Japan. Many thanks to Dr. Kouichi Nishimura who hosted the author and organised the Miyanomori ping-pong ball avalanches experiments. This work could not have been done without the help of Mr Ogura Takahiro who was responsible for filming the avalanches and transferring the film onto computer. The author would also like to thank Karstein Lied (NGI) and Mohamed Naaim (Cemagref) for providing film of snow avalanches.

References

[BAS 88] M. Basseville. Detecting changes in signals and systems, a survey. *Automatica*, 24:309–326, 1988.

[BAS 93] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, New Jersey, 1993.

[BER 94] José M. Bernardo and Adrian F. M. Smith. *Bayesian Theory*. Wiley, 1994.

[DEN 99] Y. Deng, B. S. Manjunath, and H. Shin. Color image segmentation. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1999.

[GOL 89] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. John Hopkins University Press, Baltimore, 2 edition, 1989.

[JÄH 97] Bernd Jähne. *Digital Image Processing. Concepts, Algorithms, and Scientific Applications*. Springer, Berlin, 4 edition, 1997.

[KAS 87] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active and contour models. *ijcv*, 1(4):321–331, 1987.

[LAT 97] B. Latombe, P. Ladret, F. Granada, and P. Villemain. An original active contour algorithm applied to snow avalanches. In *IPA97, 15–17 July 1997, Conference Publication No. 443*. IEE, 1997.

[LIM 90] Jae S. Lim. *Two-Dimensional Signal and Image Processing*. Prentis Hall, New Jersey, 1990.

[MA 97] W. Y. Ma and B. S. Manjunath. Edge flow: a framework of boundary detection and image segmentation. In *Proc. IEEE International Conference on Computer Vision and Pattern Recognition, San Juan, Puerto Rico, June 1997*, pages 744–749, 1997.

[MCE 01] J. N. McElwaine and K. Nishimura. Ping-pong ball avalanche experiments. *Ann. Glaciol.*, 32:241–250, 2001.

[MCE 02] J. N. McElwaine, T. Ogura, and K. Nishimura. Drag forces and ping-pong ball avalanches. To appear Seppyou (Japanese), 2002.

[NIS 98] K. Nishimura, S. Keller, J. N. McElwaine, and Y. Nohguchi. Ping-pong ball avalanche at a ski jump. *Granular Matter*, 1(2):51–56, 1998.

[O'H 99] A. O'Hagan. *Kendall's Advanced Theory of Statistics: Volume 2B - Bayesian Inference.* Arnold, London, 2 edition, 1999.

[ÓR 96] Joseph J. K. Ó Ruanaidh and William J. Fitzgerald. *Numerical Bayesian Methods Applied to Signal Processing.* Springer, 1996.

[PRA 91] William K. Pratt. *Digital Image processing.* Wiley, New York, 2 edition, 1991.