Citation for published version:

Barber, A, Brown, M, Hogbin, P & Cosker, D 2014, 'Estimating Camera Intrinsics From Motion Blur', Paper presented at Conference for Visual Media Production, London, UK United Kingdom, 13/11/14 - 14/11/14.
<https://doi.org/10.1145/2668904.2668934>

DOI:

[10.1145/2668904.2668934](https://doi.org/10.1145/2668904.2668934)

Publication date:

2014

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

© ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Barber, A, Brown, M, Hogbin, P & Cosker, D 2014, 'Estimating Camera Intrinsics From Motion Blur' Paper presented at Conference for Visual Media Production, London, United Kingdom, 13/11/14 - 14/11/14 <http://doi.acm.org/10.1145/2668904.2668934>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Estimating Camera Intrinsics From Motion Blur

Alastair Barber
University of Bath
a.e.barber@bath.ac.uk

Matthew Brown
University of Bath
m.brown@bath.ac.uk

Paul Hogbin
Double Negative
hogbin@dneg.com

Darren Cosker
University of Bath
dpc@cs.bath.ac.uk

ABSTRACT

Estimating changes in camera parameters, such as motion, focal length and exposure time over a single frame or sequence of frames is an integral part of many computer vision applications. Rapid changes in these parameters often cause motion blur to be present in an image, which can make traditional methods of feature identification and tracking difficult. Here we present a method for estimating the scale changes brought about by change in *focal length* from a single motion-blurred frame. We also use the results from two separate methods for determining the rotation of a pair of motion-blurred frames to estimate the exposure time of a frame (i.e. the *shutter angle*).

1. INTRODUCTION

Estimating motion of a camera system, both in terms of *extrinsic* (camera movement relative to the world coordinate system) and *intrinsic* camera changes (such as changes in focal length) is an important aspect of many computer vision applications. Accurate estimation of these changes throughout a film sequence is an essential part of the Visual Effects (VFX) process, as without this information, Computer Generated assets, such as characters, scenery and effects, cannot be applied convincingly to live-action footage. Often, in order to determine changes in the camera parameters, it is necessary to track individual feature points over two or more frames after filming has taken place, or use additional camera mounted hardware such as a motion capture rig, inertial measurement devices, and other devices for tracking physical changes to the lens parameters. In the case of using additional hardware, this presents challenges such as gaining acceptance on set for installation, and the additional expense of equipment and operation. There are also often many situations where such equipment would be impractical - such as outdoors or at sea, due to the reliance on additional infrastructure. However, recent developments in the field of Electromechanical Sensors has allowed for the manufacture of gyroscopes and accelerometers that are both

low cost and small. These devices are now starting to be included within cameras and can easily be mounted to them in order to provide information about their motion during filming. Examples of applications of such camera mounted devices range from assisting determining scene geometry [9] to correcting for distortions introduced by motion and camera rolling shutter [3]. One of the most significant challenges with using inertial measurement sensors to measure motion of the camera is that only changes in acceleration or rotational velocity are recorded. This can lead to significant errors in determining absolute position by integrating this data [10], and as such are rarely suitable for tracking camera motion when used alone. Devices which track physical changes in lens parameters are now commonly used in production environments and have gained acceptance across the industry - however they must be accurately synchronised to the video captured by the camera. Whilst this is now a quick process, occasionally it may not be completed correctly (if at all) for each shot, and manual alignment of the data in post-production is a time consuming and hence expensive task.

Accurate feature tracking is a reliable method of determining accurate camera motion estimations, and is an active area of research. However, there are several cases where it is difficult to get an accurate track, most noticeably when there is a fast unpredictable motion of the camera, which also often leads to a considerable amount of motion blur being present in a frame, making features undetectable. In [15], the authors present a method for determining dense optical flow in the presence of spatially-varying motion blur. This method produces good results, however calculating optical flow over an entire image can be a computationally expensive process. In [4], the authors present a method of determining in real-time and using a single motion-blurred frame, an estimate for camera rotation.

In this work, we use motion blur induced onto an image by changes in focal length and motion to track changes in two camera intrinsic parameters, namely focal length and shutter angle. We also demonstrate how, in a situation where both data from certain sensors are available, data obtained from observing blur patterns can be used to align the sensor data with the correct frame. In order to validate these results obtained from real footage, we utilise high-specification gyroscopes and a hardware focal-length tracking system. The novel contributions of this work are:

- Algorithms to track focal length and shutter angle using single images and image pairs respectively
- Accurate ground truth datasets obtained by high-end gyroscopes and hardware focal length tracking
- Analysis of the efficacy of our algorithms for intrinsic parameter estimation and sensor alignment.

In order to achieve these goals, we operate under the assumption that each frame experiences motion blur caused by either a rotation or change in focal length. One limitation that is expected in this approach is that a frame must have a sufficient amount of motion blur present, and that the accuracy of our proposed method will be influenced by the magnitude of changes in both rotation and focal length during frame exposure.

2. BACKGROUND

Using Motion blur to determine parameters of a scene is an area of current computer vision research. [7] presents a method of determining speed of a moving vehicle from a blurred image, whilst then using this information to de-blur the resulting image. Other methods, such as the one presented by Rekleitis [12] use the direction and magnitude of motion blur in the process of estimating optical flow in an image. Later work, in [15], parameterises each frame as a function of both pixel movement and motion-blur. In [15], the authors determine the derivative of the blurred frame with respect to both the motion and the blur, where the blur itself is a function of motion. Furthermore, if the exposure time is known as a fraction of the frame (*shutter angle*), the result can be further optimised. Recent work in [5] makes use of data captured from a 3D pose and position tracker attached to the camera to aid in the calculation of optical flow in images affected by motion blur. As the level of motion blur in an image is typically directly related to the exposure time of the frame, [8] and [14] use a method with a hybrid camera capturing both high and low frame-rate images of the same scene to correct images exhibiting motion blur.

Presented by Klein and Drummond in [4] is a method for determining the rotation of a camera during a single-frame exposure resulting in motion blur. In this work, the axis of rotation is derived by selecting a point through which the most normals to the edgels at a set of 'edgel' (points along an edge) points coincide. This algorithm builds on the observation that areas of motion blur will typically form edges in the image. Figure 1 shows a synthetic animation that has undergone motion blur whilst the virtual camera has been rotated, and the results of this image having undergone Canny edge detection.

In the case of the scene in figure 1, the algorithm described in [4] will estimate the centre of rotation to be at the centre of the image plane - the Z axis. In order to handle rotations around the X and Y axis, the normal line to the edge at each edgel site is expressed as the intersection of the image plane with a plane passing through the origin and an edgel site. Once the centre for rotation has been accurately determined using RANSAC (and optimised using a Levenberg-Marquardt based algorithm), the magnitude of

rotation can be determined from analysing the blur along its direction, with the intensity of pixels in the image being sampled in concentric circles centred at the estimated axis of rotation. In [4], rotation magnitude is estimated under the assumption that the blur length cannot exceed the shortest intensity ramp produced by an intensity step in the scene (i.e., the least blurred feature). Under the further assumption that the largest intensity step in each scene will span approximately the same intensity increase, the gradient of the steepest ramp to span this increase will therefore be inversely proportional to the length of the motion blur, and thus the magnitude of rotation from the camera. Their work highlights a number of important limitations in using motion blur to determine changes in camera parameters, most notably that from a single frame alone, it is not possible to determine the direction (or sign) of rotation. For this reason, it is only possible to compare the results of this algorithm with normalised values of rotation from a rate-gyroscope or other method for determining ground truth.

2.1 Intrinsic Parameters

The intrinsic parameters we consider in this work are changes in *focal length* and *shutter angle*. During image formation, rays of light reflected from the photographed subject are focused onto the film (or digital sensor) using a lens. A parameter of this lens is the *focal length* f , which determines the field-of-view, or zoom, of a scene.

If the focal length of a lens were to change whilst the sensor is exposed, it could be expected that the image will experience motion blur in a similar fashion to those described in the previous section due to changes in the field of view. An example of such an image is also shown in Fig. 1. Although the entire image has been scaled by a uniform amount, it is apparent that parts of the image are blurred by non-uniform amounts, specifically - towards the centre of the image edges will appear sharper than towards the outside. It is also clear that the 'edges' introduced by this blur converge towards the centre of the image, in a similar fashion to a translation of the camera originating from the centre of the image.

When an frame is captured, the image sensor, or film, is exposed for a short amount of time. Often, this amount of time is known and controlled by the camera operator - however there are occasions where this would be an unknown value, such as in cameras with an automatically controlled exposure. Fig. 2 shows two extracts from two video sequences of a ball falling under gravity. The left hand panel is a frame from a sequence shot with an exposure time of $1/500$ th of a second, whilst the right hand panel shows a similar scene captured with an exposure time of $1/100$ th of a second. In both frames, the ball falls at an identical speed, and in both cases the *frame rate* was set to 25 frames per second. Therefore, the left frame would be exposed for $\frac{1}{500} \div \frac{1}{25} = 0.05$ of the frame time and the right hand frame for $\frac{1}{100} \div \frac{1}{25} = 0.25$. It can be seen from Fig. 2, the frame with the longer exposure time as a fraction of the frame exhibits the largest amount of motion blur. Historically, this fraction of time for which the frame is exposed is determined by the *shutter angle*. This is so called as in cameras with mechanical shutters consisting of a rotating disk with an adjustable sector with which to expose the film, the shutter

Figure 1: Images Blurred from Camera Rotation and Focal Length Changes with Resulting Canny Edge Detection

Figure 2: Illustration of Shutter Angle and Motion Blur (25fps)

angle referred to the angle of opening of this sector. In the example from Fig. 2, the shutter angle of the second frame would be $360^\circ \times 0.25 = 90^\circ$, and a frame for which the exposure time is half the frame time would be 180° . Throughout this work, for simplicity, we refer to the values for shutter angles as fractions of the frame time.

3. METHOD

3.1 Motion from a single frame

In the case of a single motion-blurred frame undergoing rotation, we use Klein and Drummond’s original method to calculate the rotation for that frame. We also extend this method to determine a scale change brought about by a change in focal length without motion.

As shown in Fig. 1, the change in focal length (assuming the camera is not rotating or translating) adds motion blur to the image in a fashion similar to a translation towards the principal point of the image plane. Unlike the method used by Klein and Drummond to estimate for rotation, there is no need to determine the centre of the transformation as we can assume that the direction of the blur will always be towards the principal point of the image plane. Therefore, in order to determine the magnitude of blur, the intensity I , of the image along several radial lines L , is sampled from the edge of the image inwards (Fig 3). This profile is then searched for the first occurrence of an intensity step change greater than a threshold value - and the length of this change (and image position of the start and end) is recorded. This differs from the method used by Klein & Drummond, as the amount of blur induced by a change in focal length across

an image is not uniform. Therefore, edges are expected to be less blurred towards the centre of the image, and hence the shortest intensity ramp will always correspond to a minimally blurred edge towards the centre of the image. Eqn.1 describes this relationship between an image point u and the point u' after a change in focal length f : Δf .

$$\begin{aligned} u &= f \frac{X}{Z} \\ u' &= (f + \Delta f) \frac{X}{Z} \\ \frac{u'}{u} &= 1 + \frac{\Delta f}{f} \end{aligned} \quad (1)$$

Where X is a scene point of distance Z from the front nodal point of a lens.

Figure 3 shows the location of a blur region as detected by this algorithm in a synthetically blurred image, and Fig. 4 the locations of all blur regions over the image.

After a pair of points has been obtained for each radial line, a RANSAC based algorithm is used in order to determine the geometric transformation between the sets of points. In this process, the start and end points of the maximum gradient ramps from the radial search lines are represented as their respective image coordinates. The geometric transform brought about by a change in scale is then estimated to produce an estimate of the scale transform, using the points identified at each radial line. To achieve this, we adapt the standard RANSAC algorithm to take into account the observation that measuring the magnitude of motion blur by searching for the maximal gradient ramp will always produce an overestimate for the blur magnitude. This would be because even in the case where there is no blur, the sharpest edge might be several pixels in extent, and in practice, in an image with moderate motion blur, will extend several pixels beyond the blurred region. Because of this, the error metric used in the RANSAC based geometric estimation is weighted to apply a higher penalty to estimations that produce an under-estimate of the scale magnitude.

In this process, instead of maximising $\sum((r' - r)^2 < \epsilon^2)$ where r' and r are the measured and predicted radial dis-

Figure 3: A line sample location (left) and profile (right). The peak gradient has been highlighted and location marked on the image.

placements, we maximise $\sum((r' - (r + \epsilon))^2 < \epsilon^2)$. By using this method, in order to be considered an inlier, r' must be in the range r to $r + 2\epsilon$, as opposed to $r - \epsilon < r' < r + \epsilon$ as in a traditional RANSAC procedure.

This method provides an accurate estimate of the transformation between the points - whilst also rejecting outliers in the sets of points. Assuming that the motion throughout the frame is constant, and where s is the scale transform determined from the set of points, n is the per second framerate, and t is the exposure time of the frame; the scale change S brought about by the change in f can be determined by: $S = (s/t)/n$.

Figure 4: Blur length estimation along all radial lines

3.2 Motion from Multiple Frames

The optical flow of two motion-blurred images can be calculated using the baseline method described in [15]. Then, a set of feature points in the first frame are sampled using [13], and their flow vectors used to calculate corresponding points. As it is expected that there will be some outliers, we use a RANSAC algorithm similar to that described in Klein & Drummond to determine a consensus set of matching points, in order to determine rotation. Assuming a correct pair of point matches, $\hat{p}_1$ and $\hat{p}_2$, where $\hat{p} = [x, y, 1]^T$ in the image coordinate system, the line joining these points will be described as $L = \frac{\hat{p}_1 \times \hat{p}_2}{|\hat{p}_1 \times \hat{p}_2|}$. An equation for the normal to this line at its midpoint can then be deduced using the point matches and represented in the homogeneous line form. Assuming a further set of correct point matches is available, and the normal line to these can be calculated,

the point of intersection of these two normal lines should then be the centre of rotation. This estimate is selected using RANSAC. In this process, a pair of candidate points produces a centre estimate. The angle between a line drawn from this candidate and the midpoint of each other pair of point correspondences, and their respective normals is calculated - and capped at a threshold value. The point producing the lowest sum of these angles is then selected as the rotation centre. This point is then normalised, and treated as a 3D point, with the Least Mean Squares value for the angle between consensus point pairs at C being the rotation magnitude. Results obtained using this method alongside Klein and Drummond's method - using synthetic image sets are shown in the following sections. It is expected that the difference between results obtained using this method over a pair of frames, compared to the results from Klein and Drummond's method on a single frame in this pair will give an estimation of the shutter exposure time, given a known frame-rate.

In order to determine the scale change from two frames, the dense optical flow for both frames is calculated using [15]. A selection of candidate points are again selected using the method described in [13] and the corresponding points tracked from the flow vector. These points are then used in the same method as described in 3.1 in place of points from sampling along the blur, in order to produce an estimate for the change in scale of the image.

3.3 Determining Shutter Angle

By combining the results obtained from a single frame, and those from a pair of frames - it should be possible to calculate the exposure time of the frame as a fraction of the framerate, simply by dividing the motion magnitude obtained from blur by that of the frame-to-frame track. This calculation could further be simplified by using just the geometric distance between points identified by searching along the radial or circular profiles. However, it is envisaged that by performing the extra stages of rotation or scale change estimation will provide a more robust estimation for shutter angle. This is because both methods include the rejection of outliers as an important stage in the calculation of rotation.

4. RESULTS

Presented in this section are the results obtained from a variety of tests, both on synthetic and real footage. In the case of synthetic images, a single static photograph had an animated scale change applied using the Nuke compositing tool. Motion blur for this set of images was then simulated for the specified shutter opening time at each frame.

For real image sequences, an external electro-mechanical *zoom encoder* was attached to the lens on the camera used to capture the footage. This is a proprietary device that uses a geared rotary encoder meshed with the zoom ring on the lens barrel to track change in rotational position of the ring. After a simple calibration and synchronisation, this data can be used to infer the focal length at a particular frame, independently from the image captured by the camera. Such devices are commonly used throughout the visual effects and post-production process as they provide a reliable method of measuring changes in camera parameters.

For the production of ground-truth values for camera rotation, the camera was rigidly attached to a high-end rate-gyro capable of determining rotation up to a speed $175^\circ/sec$ with a standard error of $0.0005^\circ/sec/\sqrt{Hz}$. [10] presents a comprehensive description of the specifications and sources of error in inertial measurement systems.

4.1 Synthetic Tests

To test the algorithms against a synthetic and known ground truth for a change in focal length, shutter angle, and rotation, the Nuke compositing tool was used to create an animated series of frames from a single image.

4.1.1 Focal length change from a Single Frame

Results for the motion estimates for a set of rotation changes and changes in focal length are shown here. In both cases, as it is not possible to determine the direction of motion from a single frame, all of the values for both focal length change and rotation are absolute values. Fig. 7 shows a plot for results obtained for determining the change in scale induced by a change in focal length. In the left hand chart, both lines should ideally be identical, and in the right-hand scatter chart, the points should lie in an $x = y$ line. In this result, the chart on the left also shows the change in scale corrected for the known shutter exposure time of the virtual camera, which should equal the frame to frame estimate of scale (the true scale in this case). For most frames, it can be seen that the corrected estimation from blur overestimates the true scale value. This is to be expected, as if there is zero blur, the sharpest edge in the blur profile to be found (as described in Sec. 3) will still be at least one pixel (in practice on real photographs, this will likely be more) - which will therefore always result in some scale change being estimated.

4.1.2 Shutter Angle and Rotation Estimation from a pair of Frames

Figure 6 shows results from a synthetic sequence undergoing a series of varying rotations and with an animated shutter angle. Panel 1 in this figure shows the estimates for the magnitude of motion blur obtained from both the pair of frame method and single frame Klein and Drummond method, the latter being un-corrected for the known shutter exposure time. From this result it can be seen that in many cases where there is only a small amount of rotation, the single frame method from motion blur will over-estimate the amount of rotation that has occurred. However, the blur based system appears to consistently underestimate the value for rotation when there is a significant change in rotation, and this behaviour is to be expected - as detailed in Sec. 3.1, as the motion from blur will only represent a fraction of the frame time, whereas the frame to frame track will represent the full movement between frames.

Due to the noise in measuring rotation from blur, the resulting estimate for shutter angle is smoothed using a moving average filter (with a span of 4 frames) across the frame-set. This filtering is necessary because whilst the RANSAC algorithm described in Sec. 3.1 is able to reduce the effect of outlying estimates for rotation of the frame, certain conditions (further described in Sec. 5) will always produce incorrect results. The most significant source of error occurs

when the magnitude of blur in the image is not sufficient for the accurate detection of the true change in focal length or rotation. By filtering these estimates we are able to reduce the impact of these errors whilst still maintaining an acceptable level of accuracy over periods where there is only a small amount of rotation present in the frame. Furthermore, outlying estimates that predict the shutter angle to be 1 or greater (i.e. the shutter was open longer than the frame time) are also automatically discarded.

4.2 Real Footage

The algorithms described in this work were tested over a set of real images captured by a Canon 700D SLR Camera along with a 70-200mm lens. The scenes shot were indoors and in good lighting conditions. For the case of focal length estimation, a rotary encoder was attached to the lens barrel to track changes in rotation of the zoom wheel, and hence changes in the focal length. Each sequence consists of approximately 300 frames. In the case of rotation - the camera was rotated quickly and manually around an axis at various speeds and magnitudes, in order to produce a sequence that would exhibit large amounts of motion blur. Likewise, for changes in focal length, the zoom was also changed quickly and at varying speeds and magnitudes whilst filming. In all cases, the shutter speed was set to a constant 1/30th of a second - apart from the Chairs dataset where it was changed to 1/60th of a second after approximately 160 frames.

4.2.1 Shutter Angle and Rotation Estimation from a Pair of Frames

In order to validate the results produced using the 2 frame optical flow based method for determining camera rotation, the estimates obtained using this method on real footage were compared with the results obtained from a gyroscope rigidly attached to the camera during rotations around an axis. Figure 5 shows the results of this test. Ideally, the line plot for the angle estimated from optical flow against the gyroscope data should be identical, and the scatter plot for this data tend to an $x = y$ line.

Shown in figure 6 are the results obtained from rotating a camera around an axis over various magnitudes, and estimating rotation from both optical flow and blur. During shooting, the camera's shutter speed was changed from 1/30th of a second (0.83 of a frame at 25fps) to 1/60th of a second (0.415 of a frame at 25fps). Figure 6 also shows the estimated shutter angle as a fraction of the frame from the difference in estimations. As with the results from synthetic sequences, the value for shutter angle was calculated from a smoothed estimate for rotation from blur at each location above a threshold value.

4.2.2 Focal Length Change

Presented in figure 8 are the results for determining a change in focal length using a single frame using the method described previously. As with rotation from blur, the single frame method of determining focal length change is unable to determine the direction of the change, hence data from the zoom encoder (taken as the ground truth) is converted to an absolute change in value.

Rotation with Gyroscope to Validate Rotation from Optical Flow Calculation (Poster)

Comparison of Rotation from Optical Flow Calculation with Gyroscope Data (Ground Truth) - Performed on the 'Poster' Real Dataset. This dataset was produced with a rigid camera-gyroscope rig in order to validate that the estimates produced by the optical flow algorithm for rotation in the presence of motion-blur were accurate when the rotation magnitude and axis of the camera is arbitrary and otherwise unknown.

Figure 5: Comparison of Results from Optical Flow based Rotation Estimation and Gyroscope Readings

4.2.3 Alignment of Sensor Data with Video Footage

During capture of real data using both the gyroscope and zoom encoder equipment, it was necessary to synchronise the recording equipment with the video frames. This is performed by showing the camera a 'digislate' - a device which displays a time-code which refreshes at the specified framerate at the start of recording, and synchronising electronically this time-code with the data recording equipment. When the video is retrieved, the frames are manually inspected to read the time-code displayed on the device and correlate with the frame number of the sequence. Whilst this is a straightforward process to perform in a controlled environment, it is not practical in every shooting environment, e.g. if shooting from an aircraft. In such cases, manually aligning the data to the frame can be a difficult process. If an estimate can be found from frames with motion blur present as to the change in either zoom or rotation, then it could be used to assist in the alignment of the data in the case of failed synchronisation. One such way of achieving this would be the use of cross-correlation over both signals (estimate from blur and ground truth from sensors), described in detail in [11]. Shown in figure 9 are the results from using the method of focal length estimation described in this work to align data from the zoom encoder sensor, compared to the actual synchronised values. In this case, the zoom encoder started recording positions before the camera started recording frames (recording changes in zoom that were not filmed) and continued recording after the camera was stopped. Although unclear from the chart, the exact difference in manual synchronisation and estimated delay between signals was 1 frame.

Figure 9: Results for Estimating the Synchronisation between Camera and Zoom Encoder

Synthetic Sequence (Synthetic Box)

Shutter Angle and Rotation Estimates from a Synthetic Dataset (Synthetic Box Sequence)

Rotation with Changing Shutter Angle. The final two frames above have a shutter angle of half the first three. (Chairs)

Shutter Angle and Rotation Estimates from a Real Dataset ('Chairs' Sequence)

Figure 6: Results for Estimating Rotation and Shutter Angle from Blur and Optical Flow

Synthetic Focal Length Change (Zoom Synthetic Boxes)

Change in Focal Length Estimates from a Synthetic Dataset. Ideally, the dashed-blue and solid-red lines in the left-hand chart should align, and the scatter plot should tend to an $x = y$ line.

Figure 7: Results for Estimating Change in Focal Length from Blur with a Synthetic Data Set

Real Focal Length Change (Zoom Boxes)

Change in Focal Length Estimates from a Real Dataset ('Zoom Boxes' Sequence). Ideally, the green and red lines in the left-hand chart should align, and the scatter plot should tend to an $x = y$ line.

Figure 8: Results for Estimating Change in Focal Length from Blur

5. LIMITATIONS

The results obtained from using motion blur in this work do suffer from several of the limitations discussed in the original Klein & Drummond paper. Notably, one of the most significant problems encountered for the estimation of parameters using blur is the need for a significant amount of blur to be present in order to be successfully detected. This issue is also described in [4], in which the authors state that in order for the centre of rotation to be estimated accurately, motion blur of length 10 pixels or greater needs to be present in the image. In the case of determining changes in focal length from motion blur, it can be seen that the estimation of the absolute value for scale change matches more closely with the actual values at the locations of higher magnitudes of zoom change. This is also the case for rotation estimation, and so in order to estimate the shutter angle accurately - there must also be a sufficient amount of blur in the image. This problem is evidenced in figure 8 - where it can be seen in the scatter plot for estimates for focal length that in the case where the actual value for change in zoom is zero, the system will always overestimate this value. Investigating methods for determining if a frame has zero levels of blur would be a useful future work.

Another significant issue with the use of a single motion-blurred frame to estimate parameters is the inability of the system to cope with frames that have undergone more than one transformation - e.g. a rotation alongside a change in focal length.

Other limitations described in [4] for estimating parameters from blur are also present in this system, such as the intolerance to strobing, over-saturation, the requirement for pure rotation and a constant centre of rotation. However, when combined with the optical flow method described in [15], it is possible to determine the 'sign' of the rotation estimates. The method presented in [15], whilst extremely accurate (as shown by fig. 5), does have a significant limitation of requiring a large amount of resources to compute - often necessitating frames to be re-scaled prior to calculation. On average, for each blurred pair of frames at a size of 640×480 pixels, it would take approximately 30 seconds to compute an estimate for the optical flow, whereas the methods from blur would compute a result in near real-time on the same hardware (≈ 30 m/s), although this speed is highly dependent on the number of edgel sites selected and also the size of the image. Recent works in [1] and [2] have attempted to address this limitation.

Another factor that may have an effect on the result obtained for real footage would be the differences in blur introduced into a frame by a camera's rolling shutter (detailed in [6]). All of the algorithms described and used in this paper operate under the assumption that when a frame is blurred due to motion, the blur is always assumed to be constant across this frame. Due to each line of the sensor in the camera being sampled sequentially at different times, during fast movement, this assumption cannot be true. Investigating the impact and ways of minimising these effects in the algorithms using blur would be an important next stage of research.

6. CONCLUSIONS

This paper has presented a method for determining changes in focal length during a single motion blurred frame, and has produced promising results from this method that allows for the estimates to be calculated quickly. We have also extended and combined two previous works in order to estimate the shutter angle of a frame. Also presented is a method for using results obtained from these estimates to align data gathered on set by additional hardware to a video sequence. These results would be useful in a setting where it is known or suspected that one of these parameters has changed in a video sequence, however identifying the position of the change manually would be a time-consuming process.

Whilst results for determining focal length change are reasonably reliable, there is still a significant source of noise produced by combining rotation estimates from blur with frame-to-frame tracks in order to estimate the shutter angle. This is more apparent where there is a small amount of blur in the frame, and this work operates under the assumption that all frames are significantly blurred. Work to quickly identify frames that have no motion blur is ongoing. Another area of further research would be extending this system to handle frames which have been blurred by more than one type of motion - such as in the case of a translation and rotation.

7. REFERENCES

- [1] Xiaogang Chen, Jie Yang, Qiang Wu, Jiajia Zhao, and Xiangjian He. Directional high-pass filter for blurry image analysis. *Signal Processing: Image Communication*, 27(7):760 – 771, 2012.
- [2] Sunghyun Cho and Seungyong Lee. Fast motion deblurring. In *ACM SIGGRAPH Asia 2009 Papers*, SIGGRAPH Asia '09, pages 145:1–145:8, New York, NY, USA, 2009. ACM.
- [3] Alexandre Karpenko, David Jacobs, Jongmin Baek, and Marc Levoy. Digital video stabilization and rolling shutter correction using gyroscopes.
- [4] Georg Klein and Tom Drummond. A single-frame visual gyroscope. In *Proc. British Machine Vision Conference (BMVC'05)*, volume 2, pages 529–538, Oxford, September 2005. BMVA.
- [5] Wenbin Li, Yang Chen, JeeHang Lee, Gang Ren, and Darren Cosker. Robust optical flow estimation for continuous blurred scenes using rgb-motion imaging and directional filtering. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014.
- [6] Chia-Kai Liang, Li-Wen Chang, and H.H. Chen. Analysis and compensation of rolling shutter effect. *Image Processing, IEEE Transactions on*, 17(8):1323–1330, Aug 2008.
- [7] Huei-Yung Lin. Vehicle speed detection and identification from a single motion blurred image. In *Application of Computer Vision, 2005. WACV/MOTIONS '05 Volume 1. Seventh IEEE Workshops on*, volume 1, pages 461–467, Jan 2005.
- [8] S.K. Nayar and M. Ben-Ezra. Motion-based motion deblurring. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(6):689–698, June 2004.
- [9] T. Okatani and K. Deguchi. Robust estimation of

- camera translation between two images using a camera with a 3d orientation sensor. In *Pattern Recognition, 2002. Proceedings. 16th International Conference on*, volume 1, pages 275 – 278 vol.1, 2002.
- [10] Oliver J. Woodman. An introduction to inertial navigation. Technical Report UCAM-CL-TR-696, University of Cambridge, August 2007.
 - [11] S. J. Orfanidis. *Optimum Signal Processing: An Introduction. 2nd Edition*. Prentice Hall, 1996.
 - [12] Ioannis M. Rekleitis. Steerable filters and cepstral analysis for optical flow calculation from a single blurred image. In *In Vision Interface*, pages 159–166, 1996.
 - [13] Jianbo Shi and Carlo Tomasi. Good features to track. In *1994 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'94)*, pages 593 – 600, 1994.
 - [14] Yu-Wing Tai, Hao Du, M.S. Brown, and S. Lin. Image/video deblurring using a hybrid camera. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, June 2008.
 - [15] Li Zhang, T. Portz, and Hongrui Jiang. Optical flow in the presence of spatially-varying motion blur. *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 0:1752–1759, 2012.