ARTICLE IN PRESS

Computers & Graphics ■ (■■■■) ■■■–■■■

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Issue on Visual Media Production

Inferring changes in intrinsic parameters from motion blur

Alastair Barber [a,b,*], Matthew Brown [a], Paul Hogbin [b], Darren Cosker [a]

[a] Centre for Digital Entertainment, Computer Science, University of Bath, Bath, UK
[b] Double Negative Visual Effects, 160 Great Portland St., London, UK

ARTICLE INFO

Article history:
Received 16 January 2015
Received in revised form
14 May 2015
Accepted 18 May 2015

Keywords:
Visual effects
Matchmove
Feature tracking
Blur
Structure from Motion
Camera Calibration

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. In this work we describe a method for tracking changes in two camera intrinsic parameters – *shutter angle* and scale changes brought about by changes in *focal length*. We also provide a method for estimating the expected accuracy of the results obtained using these methods and evaluate how the technique performs on images with a low depth of field, and therefore likely to contain blur other than that brought about by motion.

© 2015 Elsevier Ltd. All rights reserved.

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. Commonly, the process of determining changes in camera parameters after filming is referred to as *matchmoving*. This is a process that uses structure-from-motion computer vision techniques to estimate both camera motion and 3D scene structure using corresponding feature points over multiple frames [13, p. 207]. This process can often be time-consuming, and requires the input of a skilled operator in order to produce an accurate camera track from even automatically detected and matched feature points. 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 electromechanical sensors have 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 [11] to correcting for distortions introduced by motion and camera rolling shutter [5]. 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 [12], 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

* Corresponding author at: Double Negative Visual Effects, 160 Great Portland St., London, UK.
E-mail addresses: a.e.barber@bath.ac.uk (A. Barber), m.brown@bath.ac.uk (M. Brown), paul@hogbin.org (P. Hogbin), dpc@cs.bath.ac.uk (D. Cosker).

http://dx.doi.org/10.1016/j.cag.2015.05.022
0097-8493/© 2015 Elsevier Ltd. All rights reserved.

ARTICLE IN PRESS

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. Another common method for determining camera movements is to make use of a method known as 'optical flow' across an image. In this process, a dense correspondence for each pixel across two frames is calculated. Assuming that there are a sufficient number of stationary objects in the scene, the camera's movement can be calculated using this correspondence information. Similar to automatic feature detection and matching, the process of calculating the optical flow across frames also suffers from degradation in the presence of large quantities of motion blur.

In [17], 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 [6], the authors present a method of determining in real-time and using a single motion-blurred frame, an estimate for camera rotation – using characteristics of the motion blur directly, and without selecting or matching any features from the image.

In our previous work [1], we used motion blur induced onto an image by changes in focal length and camera rotation to track changes in two camera intrinsic parameters – namely focal length and shutter angle. We used accurate hardware tracking of changes in camera parameters (the focal length change of a lens and camera rotation) to gather ground truth datasets and validate our algorithms. We also demonstrated how, in a situation where unsynchronised data from certain sensors was available alongside blurred footage, the blur patterns from frames in this footage could be used to accurately synchronise the external data with camera frames. One of the main limitations of the approach presented in [1] is that in order for an accurate estimate of focal length to be produced, there must be a sufficient amount of motion-induced blur present in the frame, along with sufficient visual texture (in this case, sharp edges). In the following sections, we give an expanded description of our method as presented in [1] for determining shutter angle and scale change brought about by focal length change. In addition to this, we present an extension to this method for validating the accuracy of such results across two new datasets in differing conditions. We also investigate the effects of a shallow depth-of field (and hence images likely to contain a significant amount of blur irrespective of motion) on both our methods.

2. Background

Our main motivation for this work is to improve the process of 'Matchmoving' for use in Visual Effects. In particular, we are interested in accurately estimating changes in camera parameters automatically and from scenes that would cause traditional structure from motion techniques based upon feature detection and matching to fail. Motion blur is often present in footage, and it is not uncommon for it to be considered a desirable artistic effect by directors in order to convey a sense of fast movement to the viewer [4]. This can often present challenges in determining an accurate camera track [13, pp. 140–143], as many current techniques for feature identification and matching rely on there being sharp corners or changes in image intensity being visible. Motion blur severely reduces the occurrence of these in an image. However, recent work has looked at using the characteristics of induced motion blur alone to determine parameters of a scene in order to avoid this limitation.

Using motion blur directly to determine parameters of a scene is an area of current computer vision research. Ref. [9] 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 [14], use the direction and magnitude of motion blur in the process of estimating optical flow in an image. Later work, in [17], parameterises each frame as a function of both pixel movement and motion-blur. In [17], 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 [7] 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, [10] and [16] 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 [6] 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. Fig. 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 Fig. 1, the algorithm described in [6] 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

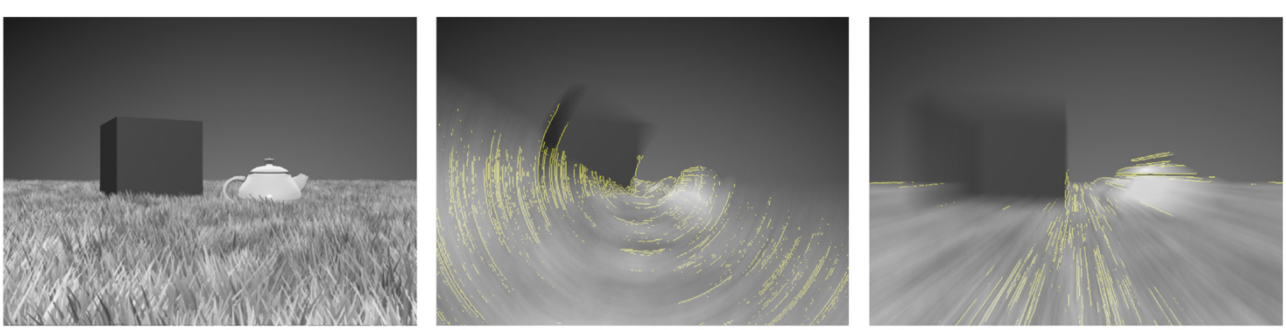

Fig. 1. Images blurred from camera rotation and focal length changes with resulting canny edge detection.

Please cite this article as: Barber A, et al. Inferring changes in intrinsic parameters from motion blur. Comput Graph (2015), http://dx.doi.org/10.1016/j.cag.2015.05.022

Download English Version:

https://daneshyari.com/en/article/6877125

Download Persian Version:

https://daneshyari.com/article/6877125

Daneshyari.com